Mind and Object.
An Essay On Intentionality

Patrik Engisch

Acknowledgments

The completion of a PhD dissertation is a rather paradoxical kind of enterprise. First, it is at once a very social and a very solitary experience. Second, one is to be held sole responsible for the final product's failures but can only be praised with caution for its better features—and there is, of course, no guarantee that it will harbor any. The two points are closely related. One writes and thinks, slowly and painfully, on one's own. Then, one discusses one's results and hypotheses with a crowd of many and is offered a turmoil of critics, feedbacks, suggestions, and new hypotheses. The overall process, at it is well-known, constitutes a strange mix of despair and exhilaration.

In my own case, I have been fortunate enough to complete this dissertation at the University of Fribourg, Switzerland, within the very stimulating and satisfying environment of the research group Experience and Reason. Among the people of EXRE, however, I owe a very special debt to my supervisor, Gianfranco Soldati. I completed pretty much all of by BA, MA, and PhD studies under his guidance and owe him a great deal of what I know not only about philosophy but also about how to do philosophy. I wish to warmly thank him for everything he did to help me throughout the years in many ways. A less ironical fellow than myself might say that he truly expressed the generosity to be expected from the pupil of a Ticinese catholic primary school. *Grazie mille per tutto questo*, Gianfranco!

Two further senior members of EXRE also deserve special thanks: Martine Nida-Rümelin and Fabian Dorsch. I learned many things from them over the years, either in class or during many philosophical discussions, and they also always have been ready to help me with more administrative issues when needed. *Martine und Fabian, herzlichen Dank für Alles!*

Fabian's passing away shortly after the submission of this dissertation was a huge shock to the EXRE community. His very kind offer to hire me as part of his new post-doc project provided me with the necessary motivation to wrap up my work in due time and I was eagerly looking forward working with him. I miss his friendship, knowledge, and leadership.

Among the other members of the EXRE community, I also owe a very special thank to Florian Wüstholz. We shared a tiny office for several years, as well as many good philosophical and less philosophical moments. Florian also took over the duty to read a great chunk of this dissertation and provided me with much philosophical comments. He also displayed his typical flair for typographical neat-picking. Florian, *ä huere grosse merci viu mau!*

I wish to extend my warm thanks to the rest of the EXRE community, past and present, for their comments, critiques, and suggestions on different things I presented throughout the years, as well as for more social occasions. I think in particular of: Emmanuel Baierlé, Miloud Belkoniene, Davor Bodrozic, Julien Bugnon, Coralie Dorsaz, Magnus Frei, Andrea Giananti, Hannes Ole Matthiessen, Jacob Naito, and Mario Schärli. Warm thanks to all!

I also had the opportunity to attend for a while the weekly seminar of the EIDOS research group in Geneva. I truly learned a lot during these sessions and also want to thank everyone down there in Geneva. On a similar note, I spent a couple of very rich months as a visiting graduate student at the LOGOS research group in Barcelona in the Spring 2013. Thanks a lot to Manuel Garcia-Carpintero, who made this visit possible, as well as to the LOGOS community as a whole—and in particular to Aurélien Darbellay—for many philosophical and less philosophical good moments.

Some of the material that found its way in this dissertation as well as bits of abandoned projects have been presented, or merely

discussed, at different occasions and several people, in addition to the one already mentioned, helped me to significantly move forward with their remarks, critiques, and suggestions. In particular, I remember helpful discussions with Johannes Brandl, Pablo Carnino, Tim Crane, Julien Deonna, Federico Lauria, Akiko Frischhut, Alex Grzankowski, Ghislain Guigon, Uriah Kriegel, Graham Peebles, Ori Simchen, Fabrice Teroni, Tobias Rosefeldt, and Alberto Voltolini.

Finally, I consider myself very lucky to have occupied a part-time teaching position while working on this dissertation. Teaching turned out to be a time-consuming activity but an intellectually and humanly highly rewarding one and I truly learned a lot from all my students of the past five years. *Un grand merci à tous!*

*

During the time of the completion of this dissertation, I was especially fortunate to lead a very rich and satisfying non-philosophical life, surrounded by loving and caring family and friends, full of these moments that made it easier to carry on. I would like to thank everyone but will limit myself to the following non-exhaustive list: my parents, Jean-Claude and Chantal; my siblings, Joris and Annick, as well as their respective families; the members of the infamous Winterchicken gang—Alex, Arnaud, David L.S., David V., François, Gil, Jean-Jacques, Nicolas, and Xavier—who truly went out of their way to shower me with the kind of good-humored silliness that is at the heart of all true friendship, as well as their respective better halves; Laurent Steiert and Barbara Horber; Beni Rohrer, Salomé Horber, and Pina; Fabiano Di Dio; Bernadette Widder; Gilles Widder and Moyli Tan; Carole Frossard and Cyprien Corminboeuf; Joane, Lucas, Roman, and Isaac Tissot-Hirsig; and the members of my two bands, Falcifer and Devil Otter Ate My Minivan.

More than anyone else, however, I would like to thank my principle source of joy and overall amazement at the greatness of life,

the intentional object of my love, Laure Widder. I dedicate this dissertation to her (as well as to Milo, Gügü, and Francis who, in their own special ways, manage to make our shared daily life even greater).

<p align="center">*</p>

On a much less joyful note, the times during which I completed this dissertation were, on a global scale, troubled ones. For the first time since World War II, Europeans have been confronted with scenes of war on their own soil. More than ever, populations in the East and, especially, the Middle-East have been the victims of men's folly, and many, including a great number of children, were forced to find a safer haven far from home. Some days, reading the news or hearing up stories from people trying to help others in distress, I truly wondered what was the point of carrying on writing a philosophy dissertation in such absurd times. In moments of doubt, I tried to find meaning in what I was doing by bringing back to memory Deleuze's claim that the goal of philosophy is to fight against stupidity. It is more than debatable whether he himself succeeded in this task, and I know this dissertation all too well to be aware that I cannot claim any special credit for myself, but to carry on trying to implement ideals of rationality as well as sensitivity to facts and evidence is the least we can do in such times.

May everyone be safe soon, wherever they wish to call it home.

Table of Contents

But with him I didn't know how to get out. I never know how to get out. One guy is battier than the other and I don't know how to get out. My trouble is I get very passionate and ecstatic. It takes me a long time to get disillusioned with the whole unreality of everything. I guess I was still loving people taking an interest in me the way he took an interest in my anti-Semitism. Yeah, he did take over from Jesus. He was going to purify me like the church. I need it black and white, I guess. There's really very little that is black and white, and I realize that the whole world is nothing but gray areas, but these mad dogmatic people, they're kind of protection, you know?

Philip Roth, *Operation Shylock*

Introduction

Perceptions, judgments, desires, hopes, fears, memories, loves, desires, and many other similar, familiar denizens of our mental life—though profoundly distinct on other counts—can all be characterized in terms of one most general feature: they all are instances of intentionality or, in other words, they all possess intentional properties.

The terms *intentional* and *intentionality* are well-known to be confusing. They are close to the word *intention*, with its familiar meaning ("I had the intention to buy chocolate but finally decided against it"), and intentions in that sense also instantiate intentional properties. The former pair, however, has a strictly technical meaning that is distinct from the familiar meaning of *intention*. Indeed, by saying that each element of the above list—and many more—are intentional, we do not mean that each is an intention in the familiar sense of the word. We mean, rather, that each of them is an instance of this capacity of the mind to be about something. Pierre Jacob clarifies the issue in his entry "Intentionality" for the *Stanford Encyclopedia of Philosophy*:

> Why is intentionality so-called? …in its philosophical usage, the meaning of the word 'intentionality' should not be confused with the ordinary meaning of the word 'intention.' As the Latin etymology of 'intentionality' indicates, the relevant idea of directedness or tension (an English word which derives from the Latin verb *tendere*) arises from pointing towards or attending to some target. (Jacob 2014)

Despite some dissident eliminativists (Quine 1960; Churchland 1981) and instrumentalists (Dennett 1987), a very large majority of

philosophers are intentional realists, *i.e.*, they endorse the thesis that the phenomenon of intentionality is real, or that intentionality is a genuine feature of reality, and not a remnant of a bygone era to be surpassed by neuroscience or a mere useful fiction. For reasons that should become clear in the upcoming pages, we shall assume intentional realism as a default position. Whatever else we human beings can be said to be, we are definitely also are creatures that must be described in intentional terms.

But what exactly *is* intentionality? We know it is the capacity of the mind to be about something. But what does "to be about something" mean? Already at this early stage, intentional realists tend to vastly disagree. Schematizing, we might divide the intentional realist group into two camps. Some intentional realists are *reductive* realists: they believe that intentionality is real but that it can, or even must, be reduced to something else—presumably to causal or causal-informational phenomena. As such, then, *to be about something* means to instantiate causal or causal-informational patterns of some sort. In the well-known words of Jerry Fodor:

> If the semantic and the intentional are real properties of things, it must be in virtue of their identity with (or maybe of their supervenience on?) properties that are themselves neither intentional nor semantic. If aboutness is real, it must be really something else. (Fodor 1987: 97)

On the other hand, some intentional realists are *non*-reductive realists, *i.e.*, they believe that intentionality is not only real but also a *fundamental* feature of reality. Hence, according to this second group, a theory of intentionality—*i.e.*, a theory of what it means for something to be about something—must appeal to irreducible intentional properties and must be given by means of an irreducibly

intentional vocabulary. As John Searle, himself a proponent of non-reductionism, puts it:

> Any attempt to reduce intentionality to something non-mental will always fail because it leaves out intentionality. (Searle 1992: 51)

What does it mean, exactly, to leave intentionality out? What exactly are these properties that are irreducibly intentional? What exactly is this intentional vocabulary? Answering these questions is a task that lies ahead of us.

There is, however, a related but distinct axis around which we can distinguish between different kinds of intentional realists: namely, the role that consciousness is supposed to play in a theory of intentionality. On the one hand, separatists believe that consciousness and intentionality are completely separable problems. On the other hand, inseparatists believe that consciousness and intentionality do not constitute separate problems but are rather, in one way or another, deeply intertwined.

Like intentionality, *consciousness* is also a confusing term, with different more or less related meanings. Its relevant meaning here is *phenomenal* consciousness, which is standardly captured in Thomas Nagel's turn of phrase, *what-it-is-likeness* (Nagel 1974). In other words, to be the bearer of a phenomenally consciousness mental state is to be the bearer of a mental state such that there is something it is like for its bearer to be in that state. By ostension, there is nothing it is like, strictly speaking, to be taller than 160 cm. There is, however, something it is like to perceive something red, to drink red wine, or to feel pain.

This second axis is related to the first one in complex ways. Indeed, ignoring separatism for the time being (we shall discuss separatism in Chapter 1, §3), we can distinguish between two kinds

of proponents of inseparatism: reductionist and non-reductionist separatists. *Reductionist* inseparatists believe that consciousness and intentionality do not constitute separate problems because there is an explanatory route from intentionality to consciousness. That is, they believe that consciousness can be reduced to intentionality and that intentionality can, in addition, be reduced to something more basic. This general idea can be implemented in different ways. One involves a higher-order theory according to which consciousness results from a two-layered structure in which a higher-order mental state is about a first-order one (Rosenthal 1986; Carruthers 2002). It can also be implemented in an intentionalist theory according to which the phenomenal properties of consciousness are accounted for in terms of the properties of representational contents (Drestke 1995; Tye 1995).

Non-reductionist inseparatists take a very different perspective on the relation between consciousness and intentionality. They believe that these two phenomena cannot be separated because we need consciousness to understand the phenomenon of intentionality. Some claim that we need consciousness to understand intentionality because not all, but at least one kind of intentional properties, *i.e.*, those characterizing mental instances of intentionality, must be explained in terms of consciousness. Another group claims that we need consciousness to understand intentionality because there is only one kind of intentionality—and it is conscious.

The debates we address in this dissertation are conceived as taking place within a rather restricted context. We shall be concerned about the debates that occur between intentional realists who are proponents of a non-reductionist kind of inseparatism and, in particular, of the thesis that intentionality is essentially conscious— or, as we shall put the claim, that a proper theory of intentionality is a theory of *intentional experiences*.

From this perspective, the debate about the nature of intentionality can be characterized as follows. We know from our first-person perspective that we are the bearers of intentional experiences. This is a piece of knowledge we all possess pre-theoretically simply by living the kind of lives we all live, *i.e.*, the lives of conscious and intentional beings. There is, however, a certain number of theoretical questions about intentional experiences. The task of a theory of intentionality is to answer them. In this dissertation, we try to answer mainly one of them: What are the concepts that a description of such experiences that take seriously the way these experiences are given to us phenomenologically as well as the way we routinely talk about them should fundamentally use?

We shall conceive of a theory of intentional experiences as a theory that aims to provide an *interpretation* of what we shall call the fundamental intentional schema (FIS). The FIS is a schema that characterizes each and every instance of intentionality: namely, "x is about y." The notion of an interpretation of that schema shall be important at two levels. First, there is much disagreement over what the little word "about" which is featured in this schema really means. Second, intentional experiences are not exhausted by being instances of that schema. On the contrary, as we shall see, they also instantiate other features. An adequate theory of intentionality should therefore be able to make the nature of the FIS intelligible and to account for these main features of intentionality. It is this task we call an interpretation of the schema.

Chapter 1, returns to consider further details of some of the main ideas introduced above. In particular, it discusses and defends the idea that a theory of intentionality should be conceived as a theory of intentional experiences. It then turns to a discussion of the FIS and to six main features of intentionality that an adequate interpretation of the FIS should account for. The next five chapters

discuss four possible interpretations of that schema. These options can be characterized in different ways. We shall essentially appeal to three different distinctions:

a) **Literal vs non-literal interpretation of the FIS**: As we shall see in more details along the way, we conceive of a literal interpretation of the FIS as an interpretation that cashes it out irreducibly in terms of the idea that for an intentional experience to be about something is for it to have something it is about or, to use a technical term, for it to have an intentional object. This shall rely on the intuition that the way intentional experiences are give to us as well as the way we talk about them speaks in favor of *fundamentally* treating intentional experiences as intentionally directed upon intentional objects.

b) **Relational vs non-relational underlying metaphysical nature of intentionality**: By "underlying metaphysical nature of intentionality" we do not mean the nature of intentional properties themselves but rather the nature of the overall phenomenon of intentionality, of which the existence of intentional properties is one central aspect. The distinction is then between theories of intentionality that do and theories of intentionality that do not in some way appeal fundamentally to the notion of a relation to account for the existence of intentional experiences.

c) **Relational vs non-relational intentional properties:** Intentional properties are those in virtue of which an intentional experience can be said to be about something

(and those for which the engineering level should purportedly be able to account). The distinction, basically, is between theories of intentionality that do and theories of intentionality that do not conceive of *about* as denoting a relation between an experience and what it is about.

In Chapter 2 and its sequel, Chapter 3, we discuss a first model of intentionality—the intentional object model (IOM)—according to which an adequate interpretation of the FIS is a literal one that conceives both of the underlying metaphysical nature of intentionality and of the nature of intentional properties as relational. As we shall see, however, the defense of these two claims requires the endorsement of a problematic theory—Meinongianism—which we discuss and ultimately reject in Chapter 3.

Chapter 4 considers a second model of intentionality—the content model (CM)—according to which an adequate interpretation of the FIS is a non-literal one that conceives of the underlying metaphysical nature of intentionality as relational but that rejects the claim that intentional properties themselves are relational.

Chapter 5 considers a third model of intentionality—adverbialism—according to which an adequate interpretation of the FIS is a non-literal one which denies both that the underlying metaphysical nature of intentionality and that the nature of intentional properties are relational.

Chapter 6, finally, considers a fourth model of intentionality—polyadism—according to which an adequate interpretation of the FIS is a literal one that denies both that the underlying metaphysical nature of intentionality and the nature of intentional properties are relational.

As such, then, the present dissertation can be regarded as an exploration into the plausibility of the perennial claim, already present in authors like Brentano and Husserl, that an adequate theory of intentional experiences should irreducibly described them in terms of the notion of intentional object. In particular, we shall try to put into perspective and highlight the importance of the following thesis, recently defended by Tim Crane (Crane 2013):

> **(T0)** A theory of intentional experiences true to the phenomenological facts must irreducibly make use of the notion of intentional object but this does not entail, contrary to the Meinongian tradition, that intentional experiences must be conceived as relations to intentional objects.

Chapter 1: Intentional Experiences

The answer, my friend, is blowin' in the wind,
The answer is blowin' in the wind.

(Bob Dylan)

1.Introduction

The overall aim of this chapter is to provide background for the upcoming chapters. It is divided into two main parts, each of which addresses a particular issue. The first issue might be put as follows: What should a theory of intentionality be a theory of? The second issue concerns what direction a theory of intentionality should take once a specific answer to the above question has been given.

With respect to the first question, we shall try to provide some reasons in favor of the claim that a theory of intentionality should be a theory of intentional experiences, *i.e.*, of conscious episodes that are about something or that have a certain subject matter. This view possesses a strong and distinguished pedigree in some philosophical circles. It is a starting point for the phenomenological tradition, inspired by Franz Brentano and founded by Edmund Husserl. It has also recently been rediscovered in contemporary analytic philosophy of mind, despite decades of debate dominated by the idea that a theory of intentionality should be couched in non-conscious terms. We believe that, like all true pieces of philosophy, it is a claim that transcends traditions.

This claim is a very substantial one. We shall not, however, try to provide a thorough and exhaustive case for it. What shall interest us in the upcoming chapters is what *follows* about the nature of

intentionality once this fundamental thesis is in place. Our aim will be first to expound the view that the proper target of a theory of intentionality are intentional experiences. We will then substantiate the view by discussing two objections and a series of replies. Finally, we turn to our second question and address the question of what shape a theory of intentionality conceived as a theory of intentional experiences should take.

The detailed plan of the chapter is the following. First, section 2 spells out what we take to be the main distinction between theories of intentionality which maintain that it is a non-conscious phenomenon and theories which maintain that it is essentially a conscious one. As we shall see, the key idea here is the one of "appearance", *i.e.*, the idea that things sensorily and cognitively consciously appear to us somehow, and that this constitutes the core of intentionality. We claim that this feature alone is sufficient to set apart intentional experiences as the proper domain of a theory of intentionality.

Second, we examine two objections, a strong one and a weak one, against the claim that a theory of intentionality should be a theory of intentional experiences. The first one rests on the general claim that consciousness cannot be intentional. This objection will be discussed in section 3. We shall see that there are good reasons to think that this argument is unwarranted. The second objection turns on the more restricted claim that an important subset of intentional states—*i.e.*, cognitive ones—are not conscious. This objection is discussed in section 4. Replies to these two objections by Charles Siewert, Galen Strawson, and Uriah Kriegel are then discussed.

Third and finally, having made our case for the claim that a theory of intentionality should be a theory of intentional experiences, we turn to the question of the general shape that a theory of intentionality should take. We argue that it should provide an

interpretation of a fundamental intentional schema that is common to all instances of intentionality, and that it should do so such that it can account for six main features of intentionality. What form such an interpretation should take is then discussed in detail in the upcoming chapters.

2. Appearances
2.1 Consciousness and Representation

What is a theory of intentionality a theory of? Well, clearly enough, it is a theory of the phenomenon of intentionality or, to put it in different words, it is a theory of the property of being intentional. However, only a superficial glimpse at the recent literature on intentionality would be required to reveal that very different things have been deemed "intentional," and, hence, that different phenomena have been taken as instances of intentionality.

Let us start by imagining a primitive creature that possesses the evolutionary selected function of representing its environment. It is, say, an immobile creature that lives on the dark bottom of an ocean which can, however, reach out for prey when it detects a presence in its vicinity. It is probably not very good at it and is not very discriminative and, hence, it sometimes grabs a piece of crap floating around or merely stabs in the dark. Still, it seems legitimate to describe this creature as being a bearer of some states that have the evolutionary selected function of representing its environment as being some way—states that have content that we may describe as *There-is-prey-over-there* content or something similar.

A very plausible theory of this capacity might take it as relying on some causal relations between the creature and its environment. These relations, admittedly, have been there from the start and might not always have been assigned the function they now have, but

evolution has produced its effect and the creature is now endowed with some kind of sensory apparatus and some kind of primitive mind, and it can, as a result, bear representational states whose function it is to detect the presence of food in the vicinity.

For instance, let us imagine that our primitive creature detects the presence of prey by means of a sensory apparatus that is sensitive to chemical changes induced by the presence of other organisms in its vicinity. Presumably, the causal-chemical relation between our primitive creature and its prey must have pre-existed the emergence of the organism's capacity to gather information about its environment *via* this causal chain. Evolution then selected this particular relation as a sufficiently reliable and effective relation for tracking the presence of food and assigned to it a particular representational function. Like all functions—and contrary to mere causal chains—however, it sometimes misfires such that the creature sometimes misrepresents its environment.

Interestingly, the situation of this primitive creature is very much like that of many artifacts engineers have been building over the ages. Take the example of a primitive speedometer. Whenever a vehicle is in motion, there is a natural relation between the speed the speedometer indicates and the number of times the wheels of the vehicle complete a full turn over a set period of time. For instance, let us stipulate that, if an unspecified vehicle V is in motion at a speed of 20 km/h, then its wheels complete three full turns per second. This, admittedly, is a causal relation that obtains *de facto* when the vehicle is in motion. Imagine next an engineer to whom has been assigned the task of creating a device that indicates the speed of the car at each moment. Quickly enough, he discovers this already existing relation between the vehicle and its wheels and sets himself to use it to complete the task he has been assigned. To do this, he realizes that he must create a device that extracts information about

the speed of *V* by using this preexisting relation between the vehicle and the number of times its wheels complete a full turn in a set period of time.

The completion of this project would give birth to a speedometer: a device that assigns to the relation obtaining *de facto* between the car and its wheels the function of gathering information about the car's speed. Note, however, that here again there is no guarantee that the speedometer will always be successful in the completion of its function. It might malfunction in extreme meteorological conditions or simply break. It would, however, still have the function that the engineer assigned to it: indicating speed.

The case of our primitive creature and that of the speedometer are of course very different. The speedometer is an artifact while the creature is a naturally evolved being. As such, the assignment of the representational function to the preexisting causal relation in the case of the primitive creature is a naturally evolved one. In the case of the speedometer, on the other hand, the function is assigned conventionally by an engineer. Still, there is a close kinship between the two cases. Both use natural relations to which functions have been assigned to represent states of the world.

Fred Dretske, one of the most important participants in the debate over the nature of intentionality in the 20ᵗʰ century, defends the idea of such a kinship between the case of our primitive creature and that of the speedometer. As he puts it in "The Nature of Thought":

> In the case of artifacts (like instruments, diagrams, gauges, and language), the functions that convert an otherwise eligible event or condition (one that *can* carry the relevant information) into a representation come from us. They are conventional in the same way all symbols are conventional. *We* determine what the function of

these artifacts will be and, hence, whether they will produce representations and, if so, of what. Since thought is not conventional, at least not in this way, the functions that convert otherwise eligible internal events and conditions in the nervous system into representations (and, hence, thoughts) are natural, not conventional. Just as they give the heart a certain blood-pumping function, evolutionary histories give various sense organs their information-providing function – thereby making our experience of the world, the product these organs produce by way of performing their function, a (sensory) representation of the world. (Dretske 2000: 234-235)

Many details about Dretske's theory should be filled in. What interests us here, however, is the following general idea it embodies: Representation is a naturally evolved phenomenon that can be found in very primitive creatures and that has a general structure—the assignment of representational functions to pre-existent causal relations—that is spread across all kinds of representational systems, both natural and artifactual, independently of their level of complexity. Accordingly, representation is a completely natural phenomenon in the sense of a phenomenon that is explicable entirely in terms of scientific categories (physical, biological, and evolutionary) without the mention of consciousness. Given this theory, the only important difference between the various kinds of representational systems is their natural or artifactual character. Other differences such as, say, the fact that some representational systems are systems of conscious creatures, are at best epiphenomenal. The idea that intentionality is something distinct from representation understood in that sense is then wrongheaded.

Such a general view of the nature of intentionality is crisply summarized by Galen Strawon in his paper, "Intentionality and Experience":

> Many present-day philosophers quickly start talking about experienceless entities like robots and pictures, computers and books, when they talk about intentionality, claiming that such things can be in intentional states or "have' intentionality even if they are not mental beings. This is extremely startling to those unfamiliar with the current debate, but the link is made as follows. First, we naturally say that such experienceless or non-mental entities are about or of things, or are in states that are about or of things. Second, it has come to seem natural to say that the problem of intentionality is nothing other than the problem of how natural phenomena can be about things or of things. Intentionality is thus equated with aboutness-or-ofness, which I will call *aboutness* for short, and the conclusion that non-mental entities can have intentionality follows immediately. (Strawson 2008: 257)

Such a general picture is seducing but, as Strawson points out, it might seem "extremely startling"—and perhaps even dubious[1]. For instance, a philosopher of a different temperament might consider Dretske's theory and be surprised that it does not really take into account the fact that some creatures are not only bearers of representational states but are also conscious creatures. And this overlooked fact, our philosopher might claim, should impact the

---

[1] Note that Strawson is making his claim in terms of a artifacts like robots and pictures, and not in terms of a primitive creature like us. However, as we stipulated that our creature is not conscious, the only relevant difference between robots and our creature is their artifactual nature, which is irrelevant for the point made here.

general theory of representation and, in particular, of intentionality in ways that people like Dretske have staunchly remained oblivious to.

Our philosopher might reason as follows. Imagine for the sake of argument that our primitive creature, though not too primitive to have low-level representational states, nonetheless is too primitive to have conscious states. That is, its sensory detectors detect the presence of prey, treat this information, and send it to some central nervous system under the form of an internal representation that encodes information which is then used by the creature to react in a certain way but, intuitively enough, the creature is conscious at no point in that process. Of course, our philosopher might not challenge the claim that this fact is sufficient to reasonably say that this creature is able to represent its environment.[2] Still, she might point out, this does not entail that this creature is such that its environment *appears* to it somehow. In other words, it does not entail that this creature undergoes experiences that are about its environment—*e.g.*, experiences of the dark and cold landscape of the deep sea or experiences of being presented with different kinds of strange fishes and sea slugs—that is, experiences that we, conscious human beings, would presumably undergo if we were to occupy its situation. And this very fact, our philosopher might claim, makes a huge difference.

According to our philosopher, hence, Dretske might be right that there is a rather well-spread natural phenomenon that we may call "representation" and that consists essentially in some evolutionarily selected systems tracking information or something along these lines. But, our philosopher might go on, this does not yet mean that all kinds of representations are properly described in these

---

[2] Referring back to the above quote, Strawson would say "'have aboutness'" but we shall rather choose to use the term "represent" for reasons that shall become obvious in the remaining of this chapter.

terms. Whenever consciousness kicks in, our philosopher might claim, something different appears: namely, experiences that are such that they make a world appear to a subject. And this, admittedly, is a much different phenomenon than that of, say, a primitive creature representing its environment in a completely non-conscious way. In fact, our philosopher could say, this could well amount to the difference between representation and intentionality proper.

2.2 Blindsight and Appearances

The empirical condition of blindsight can be helpfully used to drive this point home. Blindsighters are subjects who have suffered damage to the primary visual cortex and that, as a result, fail to visually experience anything in certain portions of their visual field. Visually confronted with a certain object—*e.g.*, a tennis ball—blindsighters report that they do not see anything. However, when asked to guess what they have been visually confronted with, their answers are correct at a rate significantly above chance and they also demonstrate some reliable level of success at interacting with objects they report not seeing—for example, catching a tennis ball they report not seeing.

The general conclusion to be drawn from cases of blindsight is that blindsighters are functionally similar to normally sighted people despite the fact that they report not seeing anything. This, admittedly, should be accepted by everyone. It is not a theoretical statement but a mere report of the behavior of people who suffer from blindsight. This said, however, one might wonder: What is the more specific conclusion to draw from cases of blindsight?

The general situation over the more specific interpretation of blindsight is well-summarized by Charles Siewert in the following passage:

To say that such subjects have "blindsight" is to say that in *one* sense they *do see* the relevant stimulus, and in *another* they *do not*. Thus in a sense they see something, and in a sense they are blind to it. How could this be so? Well, first, consider a specifically visual sense of "look" in which no object in a pitch dark room *looks* any way at all to a person. Second, interpret "see" in such a way that a person cannot be rightly said to *see* something that *looks* to her no way at all. Then, regarding blindsight, we say: in this sense, the blindsighter correctly denies *seeing* the stimulus, even though she correctly discriminates it (in verbal judgments, in movement) because retinal stimulation from it triggers activity in what's left of her visual system. So in a sense she's blind to the stimulus—it doesn't look any way to her— and in another she sees it. For the kind of discrimination she does have could also be regarded as a kind of "seeing." (Siewert 2013: 241-242)

According to Siewert, cases of blindsight can be interpreted in two different ways. That is, one can choose to individuate perceptual states either phenomenally or non-phenomenally. If one chooses to individuate perceptual states phenomenally, then blindsighters cannot be said to see anything even though they might be said to be able to gather information about their environment in some way. If, on the other hand, one chooses to individuate perceptual states non-phenomenally—*e.g.*, in pure causal-functional terms—then blindsighters can be said to see a stimulus.

Which of these two is the correct way to individuate perceptual states? As a matter of fact, we do not need at this point to side with either option. Indeed, we can remain neutral with respect to the question of the proper individuation conditions of the perceptual state while still consistently making another distinct and substantial claim about appearances: Namely, that whatever is the true claim to

make about blindsighters, it remains true that they fail to be perceptually appeared to in any way.

Now, about this difference, one could say either of two things. First, that it marks no real difference between normal and blindsighters. Or, alternatively, that it marks a substantial distinction between normal sighters and blindsighters. Namely, the first ones do not only gather information about their environment in a way that allow them to interact with it in a way that relies on this information. Additionally, they also are *appeared* to somehow and this, our philosopher might claim, marks the distinction between representation and intentionality proper.

2.3 Generalization

Can we now say something more general about the difference between non-conscious and conscious representation? A helpful and systematic distinction is offered by Colin McGinn in the following passage:

> I doubt that the self-same kind of content possessed by a conscious perceptual experience, say, could be possessed independently of consciousness; such content seems essentially conscious, shot through with subjectivity. This is because of the Janus-faced character of conscious content: it involves presence to a subject, and hence a subjective point of view. Remove the inward-looking face of [conscious] and you remove something integral… (McGinn 1988/1997: 300)

In this passage, McGinn expresses the claim that, though different kinds of things can be said to be bearers of representational states, something peculiar occurs in the case of conscious representational

states. In such cases, he says, the content of these intentional states is Janus-faced: on the one hand, it makes, as in the case of non-conscious representation, a claim as to a way the world is (*e.g, There is prey over there*); on the other hand, it also presents this way the world is, or could be, to a conscious subject. In other words, in such cases, a conscious subject is appeared to somehow.

Can we then generalize this claim to all intentional states? Our philosopher would likely continue. Indeed, he would point out that not all representational states are sensory ones and that, besides perceptions, we also are bearers of representational states of a non-sensory kind that may be regrouped under the umbrella term "thoughts" and that might also be said to be appearances. For instance, one may not only undergo a visual experience as of being presented with a tennis ball but also, *e.g.*, as of judging something about it ("It is a tennis ball"), imagining something about it ("I wonder what it smells like"), remembering something about it ("I have seen this tennis ball before"), or reasoning about it ("This tennis ball is a contingent being, it must then have a cause distinct from itself"). What about *these* kinds of representational states, our philosopher might ask? Are they also experiences as of being appeared to somehow, though cognitively rather than sensorily? In other words, should we claim that there are not only sensory intentional experiences but also cognitive ones?

At this point, our philosopher may make two complementary points: an historical one, and a systematic one. The historical point is that many philosophers have believed that there is such a thing as cognitive phenomenology.[3] The systematic point is that there is a very simple argument in favor of the existence of cognitive

---

[3] *Cf.*, for instance, the discussion of the historical roots of the debate over cognitive phenomenology in Soldati 2008 and Siewert 2011.

phenomenology: we know, from our own first-person perspective, that we undergo experiences of thinking pretty much all the time from the moment we wake up until we fall asleep—one might even point out that even falling asleep does not really suffice to free us from the grip of cognitive phenomenology, as many dreams are shot through with words and thoughts. These two claims put together, our philosopher might claim, give us a *prima facie* argument in favor of the claim that there is indeed cognitive phenomenology.

We shall return to this issue below, but let us admit for the sake of the present discussion that our philosopher is correct. What kind of consequences should we draw from this? Some philosophers believe that the consequences are far reaching. That is, the difference between non-conscious and conscious representation is so important that it is a mistake to believe that both can be analyzed under the single heading of a general theory of representation. Accordingly, we should delineate clearly between a theory of intentionality—namely, a theory of intentional experiences—and a theory of non-conscious representation. Galen Strawson, for instance, writes as follows:

> Some think it obvious that only mental entities or states or events in mental entities can be intentional or have intentionality; others are prepared to ascribe intentionality—*intentionality*, no less—to things that no ordinary person wishes to call mental. I take intentionality to be an essentially mental and indeed essentially experiential (conscious) phenomenon. This is terminologically unorthodox in present-day analytic philosophy, and I adopt it not so much because it's simply correct in the English that I speak but because I think it offers the best way to put things when trying to get a clear general view of the phenomenon of intentionality and, more broadly, the phenomenon of one thing's being about another. (Strawson 2008: 258)

If Strawson is correct, then we have a clear answer to our initial question: a theory of intentionality should be a theory of intentional experiences. Such a theory would be distinct from a theory of non-conscious representation, as it would be concerned with a particular, distinctive, set of questions such as the following:

- What features, if any, unify the class of intentional experiences?
- What are the basic concepts needed to provide an adequate description of these features?
- What differentiates and what unifies sensory and cognitive appearances?
- What exactly is phenomenally conscious in an intentional state—its content, its mode, or something additional?
- How many kinds of sensory and non-sensory phenomenology are there?
- What is the relation between the phenomenal character of an intentional state and the determination of its semantic aspects, *i.e.*, it having an intentional object and/or an intentional content?

The thesis that there are intentional experiences of both sensory and cognitive kinds and that these experiences constitute the proper object of a theory of intentionality has often but not always been objected to. This raises an important dialectical question: Who must bear the burden of proof? Is it those who agree that our philosopher is correct or those who disagree?

The answer to this question is likely to shift from context to context. In the context of early 20th century phenomenology, the burden of proof would clearly have been on the one who disagrees.

In the context of post-Rylian and post-Quinian 20ᵗʰ century analytic philosophy, the burden of proof would undoubtedly be on the shoulders of our philosopher. Nowadays, the burden of proof depends on where and for whom you voice the concerns of our philosopher.

Our take shall be rather opinionated. In the course of the next chapters, we shall take the point of view of this philosopher whose view we have been expounding in this section. That is, we shall explore the consequences of the claim that a theory of intentionality should be a theory of intentional experiences. More precisely, we should explore the question of the foundations of a theory of intentional experiences, *i.e.*, what are the most general and basic concepts one must use to describe intentional experiences?

We shall not, however, assume it wildly merely on the basis of the above *prima facie* arguments. The two next sections discuss some ways some authors have disagreed with our philosopher and try to make a slightly stronger case in her favor. Our aim, however, is not to prove our philosopher right, as this would require much more room that we shall allow to the issue. The aim, rather, is to clarify how much is loaded in the claim that a theory of intentionality is a theory of intentional experiences and, incidentally, to shift the burden of proof to the shoulders of those who deny that a theory of intentionality should be a theory of intentional experiences.

We shall focus on two particular objections against this claim: a strong and a weak one. The stronger objection amounts to the denial that there can be intentional experiences of whatever kind. The weaker objection, on the other hand, does not deny that there are intentional experiences; it claims instead that the identification of intentionality with such experiences overreaches, as not all kinds of intentional states can be said to be experiences. We start in the next section with the strong objection and then turn in §4 to the weak one.

3. There Are No Intentional Experiences
3.1 The Separatist Objection

The nature of the mind is generally considered to be constituted by the existence of two phenomena: intentionality and consciousness. Separatism is a standard view in the philosophy of mind that is realist about both of them. It endorses, however, a particular thesis about their relation. Indeed, according to separatism, intentionality and consciousness are not only jointly exhaustive; they are also mutually exclusive: whatever is phenomenal is not intentional and whatever is intentional is not phenomenal. The strong objection is an offspring of separatism.

At the core of separatism lies a particular conception of consciousness according to which conscious states are purely phenomenal, *i.e.*, there is something it is like to have them but they have no intentional features—they are silent about the way the world is, or could be. Such purely phenomenal states are called sensations or raw feels, the latter label supposedly elucidating the nature of these conscious states: they are raw, primitive sensations that absolutely lack any trace of intellectual sophistication. About them, Gilbert Ryle famously writes that,

> …from the mere inventory of the contents [the stream of consciousness] there would no possibility of deciding whether the creature that had these sensations was an animal or a human being; an idiot, lunatic, or a sane man; much less whether he was an ambitious and argumentative philologist or a slow-witted but industrious magistrates' clerk. (Ryle 1949: 195)

This constitutes a fundamentally democratic conception of consciousness, according to which all kinds of conscious being share the same feature in the same way, independently of where they are located on the scale of cognitive complexity and intelligence. This conception of consciousness is generally captured ostensively by means of pointing to some paradigmatic examples of allegedly pure phenomenal states such as feelings of pain and pleasure, sensations of color such as phenomenal redness, and moods such as elation or depression. To feel a pain, for instance, is said to be just this: to undergo an unpleasant feeling of a certain intensity that might be localized but is not, strictly speaking, about that location. The pain I now feel in my back is *in* my back but is not *about* my back. It is just a raw feeling of pain.

The claim that consciousness is essentially a matter of raw feels is a substantial thesis with important consequences. Some of the main consequences can be found on the following list:

a) Raw feels are not conceptual. To experience redness, hence, neither requires nor amounts to applying the concept "being red" to something. Of course, experiences of redness may play a role in the acquisition of the concept of redness, but in all cases they fall short of being constitutive of the concept, as not all creatures that experience redness have the concept of being red and, possibly, not all creatures that possess the concept of redness have phenomenal consciousness.

b) Raw feels do not have content. In other words, they do not make a claim about the world, *i.e.*, do not claim, veridically or not, that something is some way. As Ryle puts it:

> We can make mistakes of observation, but it is nonsense to speak of either making or avoiding mistakes in sensation; sensations can

neither be correct or incorrect, veridical or non-veridical. They are neither apprehensions nor misapprenhesions. Observing is finding out, or trying to find out, something, but having a sensation is neither finding out, nor trying to find out, nor failing to find out, anything. (Ryle 1949: 195)

Note that this second property must be distinguished from the first one insofar as some philosophers have promoted the idea of a non-conceptual form of content.

c) Raw feels cannot constitute the content of an intentional state such as belief without recourse to a process of interpretation. Indeed, since they are not conceptual and do not have content, and since beliefs require at least something content-like, there is no raw feel that could directly be put to use in an intentional state like a belief. At best, raw feels can cause someone to form a belief the content of which is the result of an interpretation of raw feels;

d) The relation that obtains between raw feels and intentional states, if any, is at best a purely contingent one of accompaniment. In other words, if the bearer of some kind of intentional states feels some way in conjunction with being the bearer of that intentional state, then the two events are at best contingently related and there is no constitutive relation that holds between them.

This last feature is central to separatism's treatment of cases of blindsight. We discussed above the question of whether the individuation conditions of perception should be functional or phenomenal. According to separatism, the answer is obvious: blindsighters fail to enjoy the characteristic phenomenology that accompanies the occurrence of perceptual states in normally sighted

persons. But this phenomenology is a mere raw feel—not an element that plays a constitutive role in the individuation of a perceptual state. Normally sighted persons and blindsighters can then be perceptually individuated along the same lines. Charles Siewert summarizes the view:

> The phenomenal character of an experience, some might say, is its "raw feel"—to have phenomenal experience is merely to have "sensations" of one sort or another; to have a mind with intentionality, or to have "mental representations"—*that* is something else altogether. (Siewert 1998: 217)

This conception of consciousness as raw feels can be used to formulate the strong objection against the claim that a theory of intentionality is essentially a theory of intentional experiences. It can be formulated under the form of a *reductio ad absurdum* along the following lines:

1) To be the bearer of an intentional state is to undergo a certain intentional experience;
2) Experiences are exhausted by their phenomenal features;
3) Phenomenal features are mere sensations or raw feels;
4) Hence, there cannot be intentional experiences;
5) Hence, to be the bearer of an intentional state cannot be to undergo a certain intentional experience.

Premise 1) is assumed for *reductio*. 2) is taken here as a definition of experiences and is then true by definition. 3), then, is without surprise the key premise of the argument. So far, however, we have merely discussed its content, not the reasons that might be advanced in its

favor. We shall mention two kinds of such reasons: phenomenological and methodological.

First, one might try to defend 3) on a phenomenological basis. That is, one might claim that it is a phenomenally obvious fact about phenomenally conscious mental states that they cannot be about anything. Pains, for instance, are not properly described as experiences in which something appears in some way. Of course, we do say things such as, "This is painful," when we inadvertently put our hands in contact with a burning hot surface. But this does not mean that our experience of pain is about something. It can either amount to a second-order commentary on the experience itself ("This experience of pain is painful") or to a demonstrative reference to the cause of the experience which, strictly speaking, is not painful at all. From then on, one might just generalize this point to all instances of consciousness. And, moreover, if doubts are raised about particular cases, one might always recur to the above-mentioned thesis of accompaniment to explain away the fact that some cases of phenomenal consciousness might seem to be intentional.

A second route is the following methodological one. True enough, consciousness is one of the most baffling features of reality. We simply do not know where it comes from or how to explain it, especially within the physicalist mindset of contemporary philosophy. Accordingly, one might see methodological virtue in endorsing a view like separatism, which allows one to investigate the phenomena of intentionality and consciousness separately and does not render a theory of intentionality hostage to the epistemic veil behind which consciousness stands. In other words, one might endorse 3) with the aim of pursuing a divide-and-conquer strategy.

One might distinguish between two different versions of this methodological strategy: a strong and a weak one. According to the

strong strategy, the isolation of consciousness constitutes a prelude to its elimination pure and simple. Indeed, the most radical way to deal with a theoretical troublemaker is simply to claim that there is no such thing in the first place. This, however, amount to a strategy stronger than that of separatism, which is realist about both intentionality and consciousness.

The weaker strategy does not regard isolation as a prelude to elimination but, rather, as prelude to the confinement of a theoretical troublemaker to a form of epiphenomenalism which holds that consciousness is real but that does not really play any substantial role. Consciousness, accordingly, would be the mental equivalent of our tailbone: a feature that is real but that we could well do without. As Charles Siewert summarizes the general take of separatism on consciousness:

> Confining differences in the character of conscious experience to a supposed realm of brute, unintelligent sensation can lead one to take a rather dim view of consciousness. For once this is totally dissociated from intentionality (and thus from intelligence, and character—or even from what makes us importantly different from reptiles), it may well begin to seem odd to say that we are conscious in this sense, and of questionable significance—so much that we may begin to doubt that we are, and seek somehow to dispose of this irksome philosophical peculiarity that has nothing to do with the main business of the mind. (Siewert 1998: 218)

We now possess both an argument against the claim that a theory of intentionality should essentially be a theory of intentional experiences and reasons for its key premise. In the next section, we shall see how one can resist separatism by arguing for the claim that, in the words

of Charles Siewert, "our lives are full of phenomenal sensory features that are also intentional features" (Siewert 1998: 219).

3.2 Siewert's Rejoinder

We might summarize the gist of Siewert's rejoinder against the separatists' objection as follows: bad phenomenological description makes for a dim view of consciousness. Indeed, according to Siewert, the main problem with premise 3) is that it widely under-describes the phenomenal character of perceptual states. The treatment that must be offered against separatism is then a phenomenological treatment, *i.e.*, we must describe more carefully the phenomenology of perceptual states to make the point that phenomenal consciousness goes at least in some instances beyond mere raw feels.

Siewert's starting point in this endeavor is the following stipulation about the phenomenal character of perceptual states:

> The phenomenal character of vision is how it seems for me for it *to look some way or other*—for example, the way it seems to me for it to look as if something has a certain color, or shape. And its seeming to me certain ways for it to look as if something is, for example, red, or yellow, or X-shaped, or O-shaped, is my possession of various phenomenal features. (Siewert 1998: 219 – italics added)

Two important remarks are in order about this quote. First, as already mentioned, it constitutes a stipulation. This stipulation, however, possesses a specific ground, as it originates in the phenomenological observation that the expression, *to look some way or other*, constitutes a much more adequate descriptive tool to report the phenomenology of vision that the vocabulary of raw feels. Accordingly, what blindsighters lack is not a shot of raw feels but, rather, an appearance

of a certain way the world looks to normally sighted persons. In other words, it is not just that their perceptual states fail to feel some way but, rather, that their environment fails to look some way to them.

This stipulation might of course be contested. But it seems hard to see how it could be contested on phenomenological grounds. What is required is an objection that is focused against following such a phenomenological stipulation. Such an objection could either be motivated by methodological remarks such as those offered above in favor of premise 3) or by a general suspicion against phenomenological descriptions. The first reason, however, looks suspiciously like separatism bootstrapping itself. The second one, on the other hand, is simply too strong.

Second, the general property of "looking some way or other" can be factored in more fine-grained look-types, *i.e.*, types of phenomenal characters associated with visual perception. These types, moreover, are rather rich and cover not only colors but, at least minimally, also things like shapes. They also differ from each other significantly. Looking red differs from looking yellow, and looking X-shaped differs from looking O-shaped[4]. In other words, these phenomenal types are distinctive. As Siewert puts it,

> The phenomenal features here invoked differ from one another, since the way it seems to me for it to look as if there's something red, differs from the way it seems to me for it to look as if there's something yellow, both of which differ in turn from the way it seems

---

[4] For dialectical reasons, Siewert limits itself to lower-order properties like shapes and colors but once in place, the view can naturally be extended to also cover higher-order properties, such as, *e.g.*, looking like a tennis-ball.

to me for it to look as if there's something X-shaped, while each of these also differs from the way it seems to me for it to look as if there's something O-shaped. To have such differing phenomenal features is to have experience with differences in its phenomenal character. (Siewert 1998: 219)

On the basis of these two quotes, we can now formulate the first premise of Siewert's rejoinder: that the phenomenal character of vision is not a matter of tokening some raw feels but, rather, of how it seems for something to look some way or other.

According to Siewert, however, if this premise is to serve as the basis of an argument against 3), it must be complemented by the further claim that this rich perceptual phenomenology is not the result of an act of interpretation. That is, one does not start with a raw phenomenology that, after a stage of interpretation, can be associated with a certain color or a certain shape. On the contrary, the fact that a certain conscious mental state has a certain phenomenal character, that it looks some way or other, is a basic fact about that mental state: one that obtains without the need of any kind of interpretation.

Here, Siewert draws a helpful distinction between perception and linguistic understanding. It is a plausible assumption about the meaning properties of sentences that they are not intrinsic properties of, *e.g.*, written marks on paper. For instance, the sentence, "Snow is white," *means* that snow is white only under the interpretation that the string of ink marks that compose the sentence, "Snow is white," actually means that snow is white. The case of perception seems rather different. Indeed, it is not the case that we first token some raw feel that is then interpreted and that ends up to be the phenomenal

character of, say, seeing a tennis ball[5]. On the contrary, whatever phenomenal character a perceptual state possesses, it possesses it immediately, *i.e.*, without the recourse to an interpretative step.

Siewert's next premise will now get us to the topic of intentionality. The question of giving a definition of intentionality is premature, but we can at least agree with Siewert that being assessable for accuracy is a sufficient condition for something to be intentional (Siewert 1998: 220). Indeed, intentional states have many different features. A central one, however, is the fact that at least some intentional states possess semantic features—that is, features in virtue of which they can be assessed for truth, accuracy, and related notions. Take, for instance, my belief that Obama wears black socks. This belief represents a condition, namely that Obama wears black socks, such that, if this condition obtains, then my belief turns out to be true.

Siewert's idea is then to apply this sufficient condition for intentionality to the phenomenal character of vision just described. Take, for instance, my perceptual state of seeing that there is a tennis ball in front of me. How should we describe its phenomenal character? According to Siewert, we shall say—minimally—that it looks to me as if there is a furry, yellow ball in front of me. Decisively, this phenomenal content of my perceptual state constitutes a certain condition: *i.e.* that there is a furry, yellow ball in front of me such that if this condition obtains, then my perceptual state turns out to be accurate. In other words, my instantiating this phenomenal character seems to be sufficient for the claim that I am the bearer of a perceptual intentional state. As Siewert puts it,

---

[5] Note that we do not imply here that the tennis ball is seen *as* a tennis ball.

First, consider some instance of its seeming to you as it does for it to look as if something is shaped and situated in a certain way, such as its seeming to you just as it does on a given occasion for it to look as if there is something X-shaped in a certain position. If it seems this way to you, then it appears to follow that it does look to you as if there is something X-shaped in a certain position. If this is right, then its seeming this way to you is a feature in virtue of which you are assessable for accuracy—that is to say, it is an intentional feature. For, from what we have said, if it seems to you as it does for it to look this way, then, if it is also the case that there is something X-shaped in a certain position, it follows that the way it looks to you is accurate… And the fulfillment of this added condition does not secure this consequence by furnishing the attributed phenomenal feature with an interpretation. If it is correct to think assessment for visual accuracy follows in this manner from its seeming to one as it does to us for it to look as if things are shaped and situated in certain ways—then we should say that *enormously many visual phenomenal features are intentional features*. For if you are normally sighted, then for most of the time you have your eyes open, you will have some intentional phenomenal feature, and which such feature you will have will change with the way shape, position, and size visually seem to you. (Siewert 1998: 221)

According to Siewert, then, the instantiation of at least some phenomenal features, such as those of vision, are sufficient to be assessable for accuracy, and this, in turn, is sufficient for the claim that at least some experiences—perceptual ones—are intentional. The strong objection against the thesis that a theory of intentionality should essentially be a theory of intentional experiences is then wrongheaded.

4. There Is No Phenomenal Character of Thought

4.1 What Conscious Thoughts?

We shall now turn to the weaker objection: namely, that there are no intentional experiences of a cognitive kind. Importantly, the endorsement of this second objection is compatible with Siewert's rejoinder. Indeed, there is no direct route from the endorsement of the claim that at least some instances of sensory phenomenology are intentional to the claim that there are cognitive intentional experiences. What is required to take this additional step is a decisive change of perspective with respect to the nature of consciousness. It is one thing to take on board the claim that consciousness is essentially a sensory phenomenon and to disagree about the extent to which the sensory can be intentional; it is another thing to take on board the claim that some instances of consciousness are non-sensory—*i.e.*, that there is, as it is standardly called, "cognitive phenomenology"—and that these are intentional.

The additional step that is required to move from the position made accessible by the endorsement of Siewert's objection against separatism to the generalized claim that intentionality is essentially a matter of intentional experiences is the rejection of what Galen Strawon calls, "the remarkable view" (Strawson 2011: 289): namely, the claim that,

> …the subject matter of phenomenology (the completely general study of the experiential character of experience) is nothing more than sense/feeling experience… (Strawson 2011: 289)

The debate over the rejection of the remarkable view is complex and cuts deep. As already mentioned, the existence of consciousness is considered the most difficult problem of philosophy of mind—

maybe of philosophy itself—and the recognition of the existence of a non-sensory kind of phenomenology is far from a neutral philosophical step within that debate. Theories of consciousness standardly come in two kinds: those that seek to explain consciousness in terms of a relation to something and those that consider that being conscious is an intrinsic property of a mental state, *i.e.*, one that it possesses not in virtue of standing in a relation to something. There is, however, a strong opposition to cognitive phenomenology from the relational side of the debate (*cf., inter alia,* Carruthers & Veillet 2011; Tye and Wright 2011). Indeed, it seems *prima facie* more difficult to find a relatum for cognitive experiences than for sensory ones—imagine, for instance, how weird it would be to claim that cognitive phenomenology is grounded in our relations to propositions. Hence, if the recognition of the existence of consciousness already is an important philosophical step, the recognition of cognitive phenomenology is a further substantial one. Hence the importance of making the view explicit in prevision of the next chapters.

Let us start at the beginning. Does it not seem strange to claim that there are no experiences of a cognitive kind? Does this not amount to denying the obvious? Are not we all the time overwhelmed by such experiences?

> Our conscious lives are not limited to perception, bodily sensation, and affect. With the possible exceptions of moments in which one is immersed in physical exercise, musical performance, or some form of meditation, the stream of consciousness is routinely punctuated by episodes of conscious thoughts. We deliberate about what to have for lunch, we remember forgotten intentions, we consider how best to begin a letter or end a lecture, and puzzle over the meaning of a

friend's remark and the implications of a newspaper headline. (Bayne and Montague 2011:1)

Not convinced? Consider, additionally, the all-too-familiar following example. Someone is lying on her bed at night, in the darkness and silence of her bedroom, troubled by a string of thoughts and unable to fall asleep. Someone in that situation would be hard pressed to deny that her thoughts are not conscious. Indeed, in such a context of sensory deprivation, what, if not conscious thoughts, could be responsible for keeping her awake? Hence, can we not just agree on the existence of intentional experiences of a cognitive kind on the basis of overwhelming first-person evidence?

Proponents of the weak objection, however, beg to differ. They do not deny that what is keeping our poor sleeper awake are instances of consciousness. But they certainly deny that these instances of consciousness are instances of cognitive phenomenology, *i.e.*, instance of thoughts being non-sensorily phenomenally conscious. We shall, however, distinguish between two strategies they can pursue to make their case.

According to a first strategy, there is no need to postulate a phenomenology proprietary of thoughts because thoughts are conscious but not *phenomenally* conscious. According to such a view, consciousness is not a unitary phenomenon; hence, it is not exhausted by phenomenal features. Phenomenal consciousness itself might be a unitary, strictly sensory, phenomenon, but that does not mean that all consciousness is phenomenal. There is, on the one hand, phenomenal consciousness and, on the other hand, non-phenomenal consciousness, and the latter is characteristic of thoughts.

The second strategy is more familiar. It echoes the strategy applied by separatism to the case of visual perception. As we saw

above, according to separatism, we can distinguish between the intentional and the phenomenal aspects of perception and treat them as entirely separate. As we shall now see, such a strategy can also be applied to the case of thoughts. Proponents of this second strategy therefore do not need to deny that thoughts lack phenomenal character. They merely have to claim that whatever phenomenal character thoughts may exhibit, it is purely sensory. It is not, they claim, the content of the thought itself that is conscious. Rather, conscious thoughts are imbued by phenomenal traits such as inner speech and visual imagery, and these are purely sensory. Accordingly, someone bothered by conscious thoughts as she is trying to fall asleep is not strictly speaking troubled by conscious thoughts. Rather, she is trouble by thoughts accompanied by some kind of sensory phenomenology. We shall discuss these two strategies in more detail in the two upcoming sub-sections.

4.2 Thoughts Are Only Access Conscious

What is the most minimal characterization one could give of the notion of conscious thoughts? One way to put the matter is the following, widely simplified, one. Imagine that the mind is a device that essentially treats information. Most of this treatment is non-conscious. A great part of it, moreover, is *bound* to remain non-conscious. For instance, one cannot verbally report some of the information treatment done by one's visual system. Not all such information treatment, however, is bound to remain non-conscious. Indeed, instances of it give rise to mental states—*e.g.*, beliefs—that can be verbally reported and used in reasoning. For instance, staring at my computer screen, I cannot report the inner workings of my perceptual system. But I can report that there is a computer in front

of me and use that proposition in reasoning such as, "There is a computer in front of me; hence, there is something in front of me."

One would be hard pressed to deny that this distinction—between some information treatment that is bound to remain non-conscious and some that does not—captures at least part of the notion of conscious thoughts. What is characteristic of it, however, and what accounts for its minimalism, is the fact that it does not mention anything like phenomenal consciousness. It merely describes what the mind is doing: namely, separating information poised for report and use in reasoning from information that is not. And, while doing this, it remains completely silent about how it feels like to have a mind. It is therefore able to capture at least part of the notion of conscious thoughts in a way that is independent of phenomenal consciousness. Thus, one may be willing to use this minimal notion of conscious thoughts to argue against the claim that a theory of intentionality should essentially be a theory of intentional experiences. Indeed, a great part of intentional states, thoughts, would in this case not be identifiable with experiences: *i.e.*, *phenomenally* conscious states.

This very minimal account just sketched of how thoughts can be conscious can be regarded as a simplified version of an idea proposed by Ned Block: namely, that thoughts are conscious in the sense of being "access-conscious"—a kind of consciousness that is not phenomenal. Intuitively expressed, the idea of access-consciousness is that some mental states have the property of being accessible to the language module of the brain or of being put in use to rationally control our behavior. For instance, if I judge that the door is closed, I can verbally report my judgement by uttering the sentence, "The door is closed." Similarly, the mental equivalent of this sentence can be used to rationally control my behavior—*e.g.*, to launch the intention of opening the door before going through it.

Block's claim is that these elements, which characterize access consciousness—being accessible to the language module of the brain or being put in use to rationally control our behavior—are sufficient to capture what we mean when we say that we are the bearers of conscious episodes of thinking.

Ned Block's original formulation of the distinction is put slightly more technically:

> A state is access-conscious… if, in virtue of one's having the state, a representation of its content is (1) inferentially promiscuous…, that is, poised for use as a premise in reasoning, (2) poised for rational control of action, and (3) poised for rational control of speech… These three conditions are together sufficient, but not all necessary. I regard (3) as not necessary (and not independent of the others), because I want to allow that non-linguistic animals, for example chimps, [access] conscious states. I see [access]-consciousness as a cluster concept, in which (3)—roughly, reportability—is the element of the cluster with the smallest weight, though (3) is often the best practical guide to [access]-consciousness. (Block 1985: 231)

If we come back to the example of someone trying to fall asleep but bothered by a string of thoughts, the proponent of the access consciousness argument would claim that this person is troubled by the fact that some mental states possess a disposition that is constantly triggered, somewhat overflowing one's mind with access-conscious thoughts. There is no need, hence, to claim that these troubling thoughts are phenomenally conscious. They do not constitute intentional experiences.

There is, however, a major problem with this proposal since, as many authors have remarked, an important point about access consciousness as conceived by Block is that it is a dispositional

property. Indeed, as Block puts it, it is sufficient to be "poised" to be access conscious. Being accessed, hence, is not necessary. One might wonder, however, how could something be conscious and not occurrent? That seems indeed stretching the word "consciousness" far beyond its legitimate application.

Moreover, in order for this proposal to work, proponents of access consciousness should be able to show that access consciousness is independent of phenomenal consciousness, *i.e.*, that something can be access-conscious independently of it being phenomenally conscious. According to Block, such a justification can be offered by pointing to states that are access conscious but not phenomenally conscious. He gives two such examples.

He regards the first one as contentious, but it is particularly clear. Imagine a robot computationally identical to a human being, endowed with a visual apparatus, a silicon brain, and a language module. Plausibly, this robot could be access-conscious while failing to be phenomenally conscious. According to Block, however, it is disputable whether such robots actually are possible, in a strong, conceptual, sense of possibility. He then proposes a less contentious example.

His second, less problematic, example amounts to the coining of a hypothetical condition called "superblindsight" (Block 1985: 233). Superblindsight is supposed to be a brain impairment similar to blindsight except that subjects can not only guess but also reliably *know* what they are visually in contact with. According to Block, patients who suffer from superblindsight are patients for whom it makes sense to distinguish between, "*just knowing* and knowing on the basis of a visual experience" (Block 1985: 233).

One might wonder whether there really are such things as superblindsighters. We need not, however, worry about this issue, as some authors, such as Uriah Kriegel, have localized the weakness of

Block's notion of access consciousness at some other place. Indeed, according to Kriegel, we have strong reasons to deny that access consciousness is independent of phenomenal consciousness. Indeed, if access-consciousness is a dispositional property, then it must have a categorical basis. And, according to Kriegel, the best candidate for such a basis is phenomenal consciousness. As such, access-consciousness and phenomenal consciousness might well be considered conceptually distinct, but the route for attributing access-consciousness independently of phenomenal consciousness is barred. As Kriegel puts it,

> When a mental state is access-conscious, it must also have a categorical property *in virtue of which* it is access-conscious. The state's disposition to be freely used in personal-level cognition cannot be a brute matter. There must be an explanation *why* the state is thus poised, an explanation appealing to non-dispositional properties that *account* for the state's disposition… The reason why a mental state is poised for the subject's free use in personal-level reasoning and action control, it is reasonable to suppose, is that the subject is already aware of it. Once the subject is aware of the state, she can freely make use of it in her deliberate reasoning and action control. Thus the state's free usage to those ends can be *explained* in terms of its subjective character or for-me-ness. (Kriegel 2009: 39)

What Kriegel refers to as "for-me-ness" is a specific property of phenomenal character, that he sees as having two components, a qualitative and a subjective one. This distinction, however, does not matter much for us here. What matters is the idea that access and phenomenal consciousness cannot be strictly independent insofar as, according to this weaker interpretation, access consciousness gets explained in terms of some aspects of phenomenal consciousness.

This, however, does not entail that all instances of phenomenal consciousness are instances of access consciousness. Accordingly, the possibility of these two kinds of consciousness being conceptually independent therefore remains open.

To sum up, a first strategy to support the strong objection appeals to a distinction between two kinds of consciousness: a phenomenal one and a non-phenomenal one—the latter of which is access consciousness. However, as qualified by Block, access consciousness is merely dispositional and there are good reasons to believe that its categorical basis is nothing but phenomenal consciousness itself. As such, then, the notion of access consciousness does not seem sufficient to explain the fact that there are conscious episodes of thinking.

4.3 Thoughts and Sensory Phenomenology

In this section, we shall turn to a second strategy to account for the existence of conscious episodes of thinking—a strategy reminiscent of Separatism's treatment of perceptual experiences discussed above. The gist of this strategy is to clearly separate between the cognitive and the phenomenal elements of a conscious episode of thinking and to claim that the latter elements can be accounted for in terms of a sensory phenomenology that accompanies the cognitive elements of the state instead of being a constitutive element of it.

The dialectic of this section is slightly different than that of the former, and we shall therefore begin by clarifying it. We first consider a standard argument in favor of cognitive phenomenology. This argument, as we shall see, takes the form of an inference to the best explanation and concludes that we are committed to cognitive phenomenology. We shall then see how one may claim that this adductive inference is unwarranted as the facts pointed to in the

argument by the proponents of cognitive phenomenology can alternatively be explained by appealing to sensory phenomenology alone. We then see how the proponents of cognitive phenomenology can react to this objection.

4.3.1 The Argument from Phenomenal Contrast

Most arguments in favor of cognitive phenomenology take the form of an inference to the best explanation. That is, a certain phenomenon is said to routinely occur and it is then claimed that the best explanation we possess for the occurrence of this phenomenon is the existence of cognitive phenomenology. Opponents then propose an alternative explanation of the facts put forward by the proponents of cognitive phenomenology and object to their conclusion. The facts that we shall discuss now concern so-called phenomenal contrasts. The general idea is the following. Two numerically distinct experiences are said to possess different overall phenomenal characters. It is then argued that the two experiences share the very same sensory phenomenology. From this it is concluded that there must be a non-sensory kind of phenomenology that accounts for their difference in phenomenal character.

Galen Strawson proposes a much-discussed argument along these lines in his book *Mental Reality* (Strawson 1994). Imagine two persons, Jacques (a monoglot Frenchman) and Jack (a monoglot Englishman) listening to the news in French and hearing the following sentence: "Bob Dylan a obtenu le Prix Nobel de littérature." According to Strawson, Jacques and Jack undergo different experiences. Jacques undergoes an experience as of understanding the propositional content that Bob Dylan won the Nobel Prize of literature while John undergoes a mere auditory experience as of listening to a string of words that he cannot

decipher. According to Strawson, however, Jacques and Jack are confronted with exactly the same auditory stimulus and then must undergo the very same sensory experience. Whatever phenomenal difference there is between them must then be accounted for in terms of a non-sensory kind of phenomenology: a cognitive phenomenology. As Strawson puts it,

> It is certainly true that Jacques's experience when listening to the news is very different from Jack's. And the difference between the two can be expressed by saying that Jacques, when exposed to the stream of sound, has what one may perfectly well call "an experience (as) of understanding" or "an understanding-experience", while Jack does not. Unlike Jack, Jacques automatically and involuntarily *takes* the sounds *as* signs, and indeed as words and sentences, that he automatically and involuntarily understands as *expressing certain propositions* and as representing reality as constituted in certain ways. As a result, Jacques's *experience* is quite different from Jack's. And the fact that Jacques understands what is said is not only the principal explanation of why this is so, it is also the principal description of the respect in which his experience differs from Jack's. (Strawson 1994: 5)

We may rephrase this point in the terminology of appearances. According to Strawson, Jacques and Jack are sensorily appeared to in the same way. If Jacques was able to bracket entirely his understanding of the language in such a way that he was able to pay attention only to the physical properties of the sounds he hears, then he would be appeared to in just the same way as Jack. But upon hearing the same sounds, Jacques, contrary to Jack, also is appeared to cognitively, *i.e.*, as he listen to the news, he is aware of the sounds he hears as having meaning and expressing a proposition. For

instance, hearing the sentence, "Bob Dylan a obtenu le Prix Nobel de littérature," makes him cognitively appeared to as if the proposition that Bob Dylan got the Nobel Prize in literature is true.

According to the opponents of cognitive phenomenology, however, contrast arguments of this kind are inconclusive. The fact brought to the table—*i.e.*, the experiential difference between Jacques and Jack—does not significantly raise the likelihood of there being non-sensory phenomenology. Indeed, according to them, it is possible to provide an alternative explanation of these facts that is just as good and that does not appeal to the existence of cognitive phenomenology.

Their main argumentative line is that the sensory phenomenology of contrast cases is under-described by proponents of cognitive phenomenology. They claim, for instance, that the claim that Jacques and Jack share the same sensory phenomenology is unwarranted. Jacques' understanding of French, for instance, might influence his perception of the sentence, "Bob Dylan a obtenu le Prix Nobel de littérature." This may occur on two different levels. First, his perception of the sentence might be altered by his understanding of French, *e.g.*, his ability to syntactically parse this sentence in the correct way might cause him to undergo a slightly different sensory experience. Second, his understanding of the sentence might be the cause of a change in his overall sensory experience: upon understanding that sentence, he might feel joy, or deception, or surprise—emotions, that admittedly, have a sensory component. According to the opponents of cognitive phenomenology, a combination of such elements is sufficient to explain the experiential difference between Jacques and Jack.

An important point in this rejoinder concerns the fact that one's understanding of a sentence might affect the phenomenal character. One might wonder whether this does not already commit

us to the recognition of the existence of cognitive phenomenology. Opponents of cognitive phenomenology, however, beg to differ. They assert that it is one thing to claim that understanding a sentence can influence one's overall phenomenology but that it is an altogether different matter to claim that understanding a sentence is a constitutive part of the overall phenomenal character of a conscious episode. One's conceptual abilities might, for instance, be crucial to shaping one's sensory experience, but this does not yet amount to the claim that concepts are phenomenally conscious. On the contrary, we might distinguish between the causal and the constitutive contribution that concepts make to phenomenal character and the opponents might object to the proponents of cognitive phenomenology that what they take to be a constitutive contribution is in fact a mere causal one.

Carruthers and Veillet write the following about the role that concepts play in visual perception:

> … there is every reason to think that concepts do their work in perception unconsciously, in interactions with incoming nonconceptual representations within the visual system before any representations become conscious… So the contribution made by concepts to one's perceptual phenomenology can be merely causal even if there isn't any time differential discernible within phenomenal experience itself. (Carruthers & Veillet 2011: 40)

The idea of this passage is the following. It is obvious that the phenomenal character of a visual experience of watching a bird is different before and after learning that this bird is, say, a bluejay. Indeed, before acquiring the capacity to visually recognize bluejays, one sees a bluejay merely as a bird with beautiful shades of blue but does not see it as a bluejay. Once one has acquired the conceptual

ability to apply the concept "bluejay," however, one may see the exact same bird as a bluejay. According to Carruthers and Veillet, however, it is a mistake to think this means that, after concept acquisition, a new layer of non-sensory phenomenal character gets added to one's overall phenomenal character, *i.e.*, a bluejay cognitive phenomenology. Rather, they claim that the concept "bluejay", once acquired, interacts at a non-conscious level with the sensory manifold of visual experience in such a way that the visual experience is differently structured, and hence feels different, before and after acquisition of the concept.

Carruthers and Veillet speak in the above quote only of visual perception, but they are confident that their model can be extended to apply to thoughts as well. In the case of Jacques and Jack, the conceptual abilities Jacques manifests in his understanding of French might be said to non-consciously structure his experience such that his sensory experience feels different than Jack's. In the words of Strawson, the opponents can generalize their claim that concepts play a merely causal role because, "sub-experiential operation[s] structure the form of our attention in certain ways" (Strawson 2011: 302).

Let us take stock. In light of these remarks, it appears that phenomenal contrast arguments such as that offered above by Strawson are not strong enough to place the burden of proof on the shoulders of the opponents of cognitive phenomenology. Indeed, it does not seem that the elements pointed by Strawson alone are strong enough to significantly raise the likelihood of there being cognitive phenomenology. We must then dig deeper to reach our aim and turn to better arguments in favor of cognitive phenomenology in the currently booming literature on the subject. We shall now turn to two such arguments: one by Uriah Kriegel and one by Galen Strawson.

As we shall see, the two arguments use different strategies. Kriegel's strategy might reasonably be called an isolation strategy, as he mounts a hypothetical case of phenomenal contrast that aims to isolate the contribution of cognitive phenomenology from sensory inputs. Strawson's strategy, on the other hand, might rather be called a sufficiency strategy. That is, he provides phenomenological descriptions of cognitive experiences that have a level of detail so overwhelming that the sensory strategy clearly appears to be insufficient. We present each strategy in turn.

4.3.2 Kriegel's Isolation Strategy

In his recent book, *The Varieties of Consciousness*, Uriah Kriegel (Kriegel 2015) proposes a thought experiment aimed to isolate the contribution cognitive phenomenology adduces to the mental life of a hypothetical subject. The argument starts with the following characterization of phenomenal properties:

> For any mental property F, F is a phenomenal property iff there is a rationally warranted appearance (to a sufficiently reflective but otherwise normal actual subject) of a distinctive (*e.g.*, empirical and nonderivative) explanatory gap between F and physical properties. (Kriegel 2015: 53)

Let us unpack this bi-conditional a bit. It is generally admitted that the existence of conscious phenomena raises concerns for physicalism, giving rise to what Kriegel calls a certain, "distinctively philosophical anxiety" (Kriegel 2015: 48). The case is well-known: we have trouble understanding how something like phenomenal consciousness could be explained within the boundaries of a physicalist theory. Pain does not seem to be reducible to a

functionalist analysis because, even though we may distinguish input conditions and behavioral outputs characteristic of pain, what constitutes the nature of pain—its being painful—eludes such a functional characterization. Similarly, how the painful character of pain could emerge from the agitation of some neurons is a fact that we simply cannot wrap our minds around. This mysterious character of consciousness can be best captured by imagining an ideal reasoner who knows all the physical facts and all the phenomenal facts. Who is then able to understand how the latter can be explained in terms of the former? Admittedly, there is something here that simply seems to elude even the most complete understanding. This creates what Joe Levine has famously called an "explanatory gap." (Levine 1983)[6].

According to Kriegel, we may distinguish between two main kinds of reactions to the mention of such an explanatory gap between the physical facts and phenomenal ones. One can either dismiss the notion as unwarranted—as, say, the prejudice of a certain philosophical tradition or a social construct or whatever. Alternatively, one can recognize it as rationally warranted, while duly noting that this does not entail the further step of accepting that there is indeed a gap in reality between physical and conscious facts.

---

[6] Note that if the coining of the concept of "explanatory gap" is rather recent, the fundamental idea is already to be found in the writings of William James more than a century ago, as he writes: "According to the assumptions of this book, thoughts accompany the brain's workings, and those thoughts are cognitive of realities. The whole relation is one in which we can only write down empirically, confessing that no glimpse of an explanation of it is yet in sight. That brains should give rise to a knowing consciousness at all, this is the one mystery which returns, no matter of what sort the consciousness and and of what sort the knowledge may be. Sensations, aware of mere qualities, involve the mystery as much as thoughts, aware of complex system, involve it." (James quoted in Kim 2006: 221)

Kriegel assumes that the explanatory gap of phenomenal consciousness is rationally warranted.

Second, an explanatory gap can either be derivative or non-derivative. As an example of a derivative explanatory gap, Kriegel mentions the problem of free will. Phyiscalism is incompatible with free will but it clashes with the best evidence we have for libertarianism, namely experiences as of acting freely. There is then an explanatory gap between physicalism and free will. According to Kriegel, however, in such a case the existence of an explanatory is merely derivative on the existence of consciousness. The existence of consciousness itself, on the other hand, creates a non-derivative explanatory gap.

Third, the explanatory gap of consciousness is an empirical one, *i.e.*, it concerns what we may call the concrete fabric of the world. It is an explanatory gap about properties and states of concrete entities: us. As such, it can be distinguished from other explanatory gaps, such as that between the physical and the normative, as the normative does not seem to be a feature of the physical but, in Kriegel words, of something that lies "beyond it" (Kriegel 2015: 51).

His characterization of phenomenal properties in place, Kriegel mounts a thought experiment that aims to make explicit that there is a rationally warranted appearance of a distinctive non-derivative explanatory gap between *cognitive* phenomenal properties and physical facts. The thought experiment comes in two parts. In the first part, we are asked to imagine a series of partial zombies, *i.e.*, subjects that lack some particular phenomenology. We are asked to imagine different kinds of perceptual zombies—*e.g.*, visual or auditory—who lack the distinctive phenomenology of the perceptual modality at stake in each case, while still conserving the ability to

track information about their environment. These zombies can, *mutatis mutandis*, be modeled on actual cases of blindsight.

Second, we are asked to imagine a different kind of partial zombie: namely, an algedonic one—*i.e.*, a subject that lacks any kind of phenomenology associated with pleasure and pain. Here again we possess an empirical case upon which imagination can get its grip: namely, congenital analgesia, a condition that is characterized by a congenital incapacity to feel pain. The case can then imaginatively be extended to pleasure.

Third, we are asked to imagine an emotional zombie: that is, a subject who lacks any kind of phenomenology characteristic of emotional episodes such as fear, anger, or love. Extreme forms of autism can serve here as the imaginative basis for such a zombie. As Kriegel points out, some people who suffer from autism report experiencing only a restricted set of emotions. We can then extend the case to a subject who is congenitally lacking in any kind of emotion.

These different kinds of partial zombies are supposed to cover the full extent of sensory phenomenology: perceptual, algedonic, and emotional phenomenology. On the basis of the imagination of these three partial zombies, Kriegel then asks us to operate an imaginative synthesis of all these different partial zombies to imagine a sensory zombie, *i.e.*, a subject who would lack all kinds of sensory phenomenology. As he puts it,

> Having now conceived in separation partial zombies with sensory, algedonic, and emotional deficits, or rather *lacunas*, we may perform another act of imaginative synthesis and envisage a person lacking all these phenomenologies at once. This person enjoys no sensory phenomenology, no algedonic phenomenology, and no emotional phenomenology. (Kriegel 2015: 55)

Zoe is the name of the subject who gets imagined as the result of this imaginative synthesis. Kriegel then goes on to fill up some crucial details about her. In particular, she is a mathematical genius who spends her days trying to prove mathematical propositions from a given set of axioms. The substantial question about Zoe is this: How should we describe her mental life?

According to Kriegel, despite the fact that Zoe lacks any kind of sensory phenomenology, we should still be able to say things about her:

> Often she struggles to find the solution of some problem—she feels stuck, if you will. But sometimes a nice thing happens next: suddenly "the coin drops" and she can *see*, so to speak, how the solution must go. Often on those occasions, a sudden intellectual gestalt shift makes Zoe realize what the missing element is, which results in a sort of affectively neutral upheaval of thought – a greater vivacity in her thinking. These victorious moments are very distinctive, and Zoe remembers many of them. Thus Zoe's mental life has its own inner rhythm, with new beginnings, stretches of inner flow, slowed down by occasional struggling and feeling stuck, often eventually punctured by breakthroughs of sudden insight and then starting over with a new mathematical problem. At the same time, this life is exhausted by intellectual or *cognitive* activities: thinking, considering, judging, realizing, intuiting, remembering, and so on. Importantly, they do *not* involve phenomenal experiences of satisfaction or frustration, though *some* underlying satisfaction- and frustration-characteristic processing does take place. (Kriegel 2015: 56)

The point of this description of Zoe's mental life is to show that, though she entirely lacks any kind of sensory phenomenology, we

can still imagine her undergoing phenomenal changes. Her mental life might be greatly impoverished in comparison to that of normal subjects, but she still undergoes phenomenal episodes, though not of a sensory kind. As Kriegel puts it, her mental life possesses its "own inner rhythm", as there are moments when she feels (non-sensorily) stuck and moments when there is a breakthrough as, for instance, when she realizes what should be her next inferential move in a certain proof. And since, by stipulation, she does not undergo any kind of sensory experiences, these experiences must be ones with a proprietary, non-sensory kind of phenomenology.

In other words, Kriegel's claim is that even if there is no room in such a scenario for a rationally warranted appearance of an explanatory gap between some sensory phenomenal state and some physical state, there is room for a rationally warranted appearance of an explanatory gap between a cognitive phenomenal state and some physical state. Indeed, how could we make sense of the kind of experiential breakthroughs just mentioned that Zoe sometimes goes through solely in terms of some neurons firing in her brains? There is here, as in the standard cases of instances of sensory phenomenology, an explanatory gap as to how the former could be nothing more than the latter. As Kriegel puts it,

> Having imagined Zoe's mental life, we may ask ourselves whether there is a rational appearance of the right explanatory gap for it. It seems that there obviously is. Consider an episode of sudden realization of how a proof must go. It is entirely natural to be deeply puzzled about how this episode could just *be* nothing but the vibration of so many neurons inside the darkness of the skull. (Kriegel 2015: 58)

This argument is, of course, not without difficulties. In particular, one may well doubt whether a subject like Zoe is really conceivable. Issues about conceivability are, of course, very complex and we cannot enter here into a discussion of modal epistemology. We may, however, mention two points. The first is that Zoe does not seem to constitute a logical impossibility. That is, one may agree that Zoe's case is difficult to imagine, but its conception does not seem to infringe any logical laws. The second point is that Zoe's argument seems to sensibly shift the burden of proof to the shoulders of the opponents of cognitive phenomenology. Indeed, they must now provide us with a good explanation of why a case such as Zoe's should be metaphysically impossible. As such, then, we can conclude that Kriegel's Zoe argument is successful for the general aim pursued in this chapter.

4.3.3 Strawson's Sufficiency Strategy

Let us now turn to Strawson's stronger argument. It starts with the simple observation that thoughts are often "extraordinarily interesting" (Strawson 2011: 299). Strawson, however, puts a specific gloss on this claim, *i.e.*, one *experiences* some thoughts as more or less interesting. The thought that one should always wear black socks, for instance, is experienced as completely dull. The thought that drugs—in particular meta-amphetamines—played an important and often overlooked role in the acts of barbarism committed by the Waffen SS in the Second World War is interesting. The first example does not sharpen our attention, and we quickly move on to other matters after having grasped its content. The latter example, on the other hand, strikes us an intrinsically interesting and important; it challenges our understanding of what happened during the Second World War and of the conditions that made such horrendous and despicable events

possible such that, having grasped its propositional content, we become focussed on it, on its plausibility, importance, and consequences, and quickly get lost in the intrinsic interestingness of such an episode of thinking, pondering over it for long minutes.

This is not a wild assertion or a stipulation. This simply is a plain fact about the way we experience our life as thinking creatures. According to Strawson, it can be used to mount a challenge for the opponent to cognitive phenomenology. Indeed, according to these opponents, the claim that some thoughts are experienced as extraordinarily interesting must be explained in terms that appeal to sensory phenomenal features alone. As Strawson puts it,

> If one wants to give anything like a full account of the experiential or lived character of our experience in merely sense/feeling terms, one has to be able to be able to explain, in those terms alone, how the experience of looking at a piece of paper with a few marks on it, or of hearing three small sounds, can male someone collapse in dead faint. (Strawson 2011: 299)

According to Strawson, however, such a project cannot be successful. It can be demonstrated that any kind of sensory phenomenology will always fall short of being sufficient to account for the fact that thoughts are often experienced as interesting. Cognitive phenomenology must be taken on board to explain this fact. Here is why.

Let us admit for the sake of argument that the opponent's view is correct: all phenomenology, including that of finding a certain thought interesting, is sensory phenomenology. How should one then describe a specific instance of finding a thought interesting? Presumably, one non-consciously understands the content of a certain thought and, soon afterward, tokens a certain feeling of

interest. The content of the thought, presumably, causes the feeling of interest but is not constitutive of it. The content is not phenomenally conscious; the feeling it causes is.

Sensory phenomenal characters are standardly rather rich. In the realm of the phenomenal character of colors, for instance, we can unproblematically distinguish a certain determinable-determinate structure in which the most general determinable might be "phenomenal character of color" and from which we can then distinguish a huge variety of determinates such as, *e.g.*, "phenomenal red," "phenomenal blue," or "phenomenal indigo." This accounts for the important variety of color experiences we can undergo. And, moreover, the explanation of why there is such a variety of color experiences is an easy one: the grounds of their distinctions are sensory ones. Indigo is sensorily different from purple red, and this explains their experiential difference.

The case of the sensory feeling of interest, on the other hand, must be different. On the one hand, it might also have some second-order properties. It could, for instance, be more or less intense and this, presumably, would explain the difference between the fact that some thoughts are experienced as interesting and some as extraordinarily interesting. It could also be more or less central to one's particular episode. For instance, one might be thinking about contemporary Russia while eating a cheese fondue and have one's overall phenomenal experience be centrally focused on the wonderful taste of the cheese and more peripherally centered on the content of one's thoughts. On the other hand, however, the feeling of interest would be dramatically poor in comparison with other sensory phenomenal characteristics. Indeed, that feeling is bound to remain "generic" or "monotonic" (Strawson 2011: 300) because we cannot distinguish such a rich structure in its sensory basis, *i.e.*, either marks on papers or sounds.

The problem for such a feeling of interest, however, is that there is a huge variety in how we experience our thoughts as interesting. Our cognitive life does not amount to a graph where a curve of generic interest goes up and down as our mind processes different kinds of non-conscious thoughts. On the contrary, we experience our thoughts as interesting in a huge variety of ways.

Now, truly enough, we *could* distinguish between different ways for thoughts to be interesting and individuate this feeling of interest more finely. This, however, could be done *only* by appealing to the content of thoughts as being constitutive of their phenomenal character—and this simply amounts to a recognition of cognitive phenomenology. In sum, then, we can conclude that sensory phenomenology is not sufficient to account for the fact that some thoughts are extraordinarily interesting.

5. Intermediate Conclusion

The main aim of the first part of this chapter was to introduce the idea that a theory of intentionality should be a theory of intentional experiences. We proceeded in three steps. First, we made an intuitive case for the claim, trying to highlight the fact that there is an important different between two kinds of representations: those that are conscious and those that are not. This difference was claimed to be important enough to provide us with a rationale for reserving the word *intentionality* for this sub-kind of representations that are conscious. Conscious representations are appearances, both of a sensory and of a cognitive kind, and this sets them apart from non-conscious ones. We might then reserve the word *intentionality* for a theory of intentional experiences.

Second, we considered a general, strong, objection against this fundamental claim. The gist of the objection is that the phenomenal

cannot be intentional and, hence, that there cannot be intentional *experiences* of any kind whatsoever. We then presented an argument by Charles Siewert against this claim and concluded that this objection is inconclusive.

Third, we considered a less general, weaker, objection against our fundamental claim. The gist of this objection is that there cannot be intentional experiences of a cognitive kind. We considered three rejoinders to this objection. The first one is a simple phenomenal-contrast argument. We determined that it is likely to be too weak for our aim. We then turned to two recent arguments and concluded that they are strong enough to place the burden on the shoulders of the opponent to cognitive phenomenology.

From this we can conclude the following. The claim that a theory of intentionality should be a theory of intentional experiences requires that there are intentional experiences of both a sensory and a cognitive kind. As such, it is a rather strong claim. As we saw, however, there is much in its favor. In any case, the above pages should have been sufficient to put the burden of proof on the shoulders of those who would deny that this is what a theory of intentionality should be about.

In the next section, we turn to another topic: granted that a theory of intentionality should be a theory of intentional experiences, what general shape should a theory of intentionality take? We argue that it should aim at the interpretation of a general intentional schema that characterizes all intentional experiences. This interpretation, moreover, must be such that it accommodates six main features of intentionality.

6. Intentionality and the Fundamental Intentional Schema

As we learned from the first part of this chapter, a theory of intentionality should be conceived as a theory of intentional experiences. The question we shall now tackle is how to conceive of the problem of intentionality when intentionality is conceived as essentially an experiential matter.

The plan of the section is the following. In a first sub-section, we distinguish between two levels at which one can pursue the discussion over a theory of intentionality—a constitutive and an engineering level—and make clear that we shall position ourselves at the constitutive one. A second sub-section introduces what we call the fundamental intentional schema (FIS). Its aim is to capture the greatest common factor in all kinds of intentional experience. We then introduce the idea that a theory of intentionality *qua* theory of intentional experiences should aim to provide an interpretation of this schema. In a third sub-section, finally, we discuss six properties of intentional states that an adequate interpretation of the FIS should be able to account for.

6.1 Engineering vs Constitutive Levels

A theory of intentionality can be pursued at two different levels: a "constitutive" level, on the one hand, and an "engineering" level, on the other—to use labels proposed by Michael Luntley (Luntley 1999: 7–8). The constitutive level can be characterized   as a theory of intentional experiences as such. That is, it is an account of how we should describe them, of the minimal set of concepts required to describe them, and of the different relations that exist between these concepts. It is, then, a blend of phenomenological and conceptual reflections.

The engineering level is rather different. Indeed, once the constitutive task has been completed, one might well wonder how a theory of intentionality as characterized at the constitutive level is engineered, *i.e.*, how a subject that has a body and a mind ends up being an intentional subject? In recent history of philosophy, the engineering level has essentially been identified with the task of providing a naturalistic account of intentionality, *i.e.*, an account of how the fundamental concepts of intentionality can be implemented within the causal order of the world. The account inspired by Dretske, which we glossed over at the beginning of this chapter, constitutes, among other candidates, an attempt to flawlessly integrate intentionality within the causal order.

This dissertation, however, is concerned only with the constitutive level. The main reason is the following. There is no good reason to get into the engineering task before one has a good grasp of the constitutive task. And, as we shall see, there are reasons to think that we have not yet reached a sufficient understanding of it. This fact shall enjoin us to formulate a principle that presides over the relation between the constitutive task and the engineering task. We shall call it, in a distinctly Helvetic fashion, the Principle of One-Sided Neutrality. The gist of this principle is that we should refrain from importing considerations that pertain to the engineering task within the constitutive task. In other words, we should not make the constitutive shape of a theory of intentionality such that it makes the completion of the engineering task easier to complete.

There are two related reasons that speak in favor of the endorsement of this principle. First, it is a sound policy to wait until the constitutive task has been fully completed before we turn to the engineering task. Second, it is already to presuppose much about the nature of intentionality to claim that the engineering task can be completed in a way that would suit the taste of contemporary

naturalism. It is not, however, that we oppose contemporary naturalism. It is, rather, that we take it to be wiser to complete the constitutive task in a way that remains neutral with respect to the general issue of whether intentionality is a naturalizable phenomenon.

In the next sub-section, we turn to consider what we take should be the core of a constitutive theory of intentionality *qua* theory of intentional experiences, what we call the *Fundamental Intentional Schema.*

6.2 The Fundamental Schema

From where should we start our inquiry into the constitutive nature of intentionality? We suggest that our starting point should be with what we call the Fundamental Intentional Schema (FIS). What do we mean by this? We mean that, at an acceptable level of generality, we can group together all instances of intentionality as instances of a single, fundamental, schema, and that to give a philosophical theory of intentionality is essentially to understand the philosophical intricacies that lie behind this schema. Let us unpack these two claims.

The fundamental schema is the admittedly not very impressive following one:

Fundamental Intentional Schema (FIS): x is about y.

To begin with, this schema is a schema and, as such, it is not supposed to mean anything on its own. Rather, it is taken to be an abstraction of instances of intentionality that aims to capture what they all have in common at a fundamental level. How do we arrive at it? Well, is it not an obvious, phenomenologically accessible fact that

all intentional experiences are about something? Is it not an obvious fact that they all have what Tim Crane calls, "an object… a subject matter… a topic" (Crane 2013: 7)? Intentional experiences are essentially those experiences through which we take a perspective on the world—*i.e.*, through which we think, judge, perceive, and so on. All these experiences are about something—this is what they all have in common, their most fundamental commonality.

This fact, one might observe, is directly reflected in the surface grammar of intentional sentences such as, "I am thinking about Obama". That said, one might also observe that not all intentional sentences do share such a surface grammar. Indeed, a great deal of intentional states are described by means of sentences formed out of a that-clause, such as, "I resent the fact that Obama will soon be out of office". Grammar, however, is not the judge we must resort to. Indeed, despite the fact that a sentence like the latter does not feature a surface grammar that directly mirrors the fundamental intentional schema, there is no reason to doubt that intentional states reported by such sentences can also be taken as instances of the fundamental schema, for it is clear that the thought that Obama will soon be gone is about something: it has a subject matter or a topic. Hence, there is little doubt that it is an instance of something that is about something else. What exactly is it about, one might ask? Obama, or the fact that Obama will soon be out of office are the answers that come to mind. Nevertheless, we shall have to wait until Chapter 2 and the clarification of the notion of intentional object to answer this question in a satisfactory manner.

At this point, one might wonder why it is precisely this schema that we take as fundamental or, alternatively, why we choose to formulate it in these terms. Why not choose an alternative formulation, such as, *e.g.*, "x is presented with y"? As we see the matter, the FIS is formulated in these terms to assure us a maximal

level of generality. Indeed, some philosophers distinguish between kinds of intentional states such as, for instance, presentations and representations. Everyone should agree, however, that all kinds of intentional states have something in common. Our suggestion is that they all are instances of the FIS.

The fundamental intentional schema looks deceptively simple. This surface simplicity, however, hides the philosophical perplexities it contains. The FIS might reveal something about the fundamental nature of intentionality, but it certainly does not exhaust it. Endorsement of the FIS is not endorsement of a theory of intentionality. It is, rather, endorsement of intentionality *qua* philosophical problem. The task of a theory of intentionality should hence be to understand what exactly it means for an intentional experience to be about something and what the conceptual commitments that applications of the FIS carry with them are, either explicitly or implicitly. The rest of this chapter points to what we take to be six main features of intentionality.

6.3 Six Main Features of Intentionality

The following six features quickly emerge from reflection on the fundamental intentional schema and its different instances. The claim, however, is neither that these six features exhaust the constitutive nature of intentionality nor that all instances of intentionality instantiate all of them. It is, rather, that a theory of intentionality that fails to adequately capture all of them will turn out to be inadequate. The constitutive task of a theory of intentionality can then be understood as the task of providing a philosophical interpretation of the FIS such that it can account for these six main features.

As we shall see, however, different theories of intentionality might account for these features in different ways. Some features

might be accounted for merely derivatively—that is, as following from some more basic features of intentionality. Some other features, on the other hand, might be considered fundamental. In introducing the features, however, we make no claim as to their derivative or non-derivative nature. These issues shall be addressed in the next chapters.

6.3.1 Aboutness

The first of these features we shall call "aboutness". It concerns what we may call the essence of intentionality: namely, the fact that intentional states are about things. One's belief that Francis is a sheep, for instance, is a belief that is about Francis. This seems easy enough, but, as we shall see in later chapters, there is a great deal of room for philosophical hair-splitting as to what we mean exactly by "being about something."

At this point, however, we shall merely point to the strangeness of aboutness itself. Let us call the instantiation of a property by something a *state*. Some states, however, are more puzzling than others, and intentional states—namely, the instantiation by an intentional being of intentional properties—have a quite unique and fantastic feature: they are not only constituents of the world but are also about it. This makes intentional states quite a fascinating topic. What, exactly, separates them from the rest of the non-intentional states? What are their most fundamental features? These are some of the fundamental questions raised by the existence of aboutness as the central feature of intentional state.

The pervasiveness of the slogan, "intentionality is aboutness," is in this respect deceitful. Indeed, to claim that intentionality is aboutness merely amounts to naming the problem. What must be understood is the meaning of the little word *about*, and, consequently, what exactly it means for an intentional experience to be about

something. In this sense, the claim that all instances of intentionality are instances of the FIS is not an explanation of intentionality. It merely points to what must be explained. And aboutness is the most fundamental of all the features of intentionality that must be accounted for.

6.3.2 Aspectuality

Imagine you are standing in front of a statue in a museum. Eyes opened, you are enjoying a visual experience as of being presented with that statue. Note, however, that though you enjoy a visual experience as of being presented with a three-dimensional object, you do not have visual access to the object in its entirety. Rather, you have visual access only to those aspects of the statue that are present in your perceptual field. The back of the statue, for instance, remains hidden to your senses.

This simple example illustrates the following idea: sensory access to an object by means of perceptual intentionality is always access to aspects of an object and never to an object in its entirety. Perceptual intentionality, in other words, always occurs from a certain point of view: namely, the point of view, to be characterized spatiotemporally, that your eyes occupy at a given time. That perceptual intentionality always occurs from a certain point of view and that what can be experienced from that point of view always depends on that point of view is an instance of what we shall call *aspectuality*.

Aspectuality, however, though it is probably most easily exemplified in the case of perception, is not unique to perceptual intentionality. It is, rather, a feature that pervades all instances of intentionality. Take the following simple example: John believes that Phosphorus is a star. This belief is about something, namely

Phosphorus. Now, as a matter of facts, Phosphorus and Hesperus are one and the same thing. Does this entail that John also believes that Hesperus is a star? No, granted that John is ignorant of the fact that Hesperus and Phosphorus are one and the same thing, there is no inference that goes from one belief to the other.

What is this example supposed to show? Since Phosphorus and Hesperus are one and the same thing, then if John's belief is about Phosphorus it should be also about Hesperus. But the fact that the above-mentioned inference does not go through speaks against this claim. In fact, what the failure of this inference shows us is that John's belief is not about Phosphorus *simpliciter*. It is about Phosphorus *qua* Phosphorus and, for instance, not about Phosphorus *qua* Hesperus. That is, John's belief about Phosphorus indeed is about Phosphorus, but it is about it only under a certain aspect: what we may call the "Phosphorus" aspect, to the exclusion of other aspects, such as, for instance the "Hesperus" aspect.

The situation here seems to be completely parallel to the perceptual case. When one thinks about something, one must be bound to think about that thing only under a certain aspect, *e.g.*, as Phosphorus or as Hesperus. We cannot, of course, speak here of points of view in the literal sense used in the perceptual case. That one is able to think about something only under a certain aspect is hence not always to be understood as a function of one's location in space and time. The idea, rather, is that just as there are physical points of view, there also are purely mental points of view from which one directs thoughts at the world.

With the introduction of aspectuality we introduce a particular kind of problem for the FIS. Indeed, aspectuality, unlike aboutness, cannot be read off the FIS. Rather, it surfaces upon analysis of instances of it. This speaks in favor of the FIS needing unfolding, or interpretation, to get a fuller understanding of intentionality. What, in

the light of aspectuality, this unfolding should look like is a task for a constitutive account of intentionality.

6.3.3 Semantic Normativity

Intentional states can be said to be normative in more than one sense. There are, for instance, normative relations that obtain between intentional states. A perceptual state as of being presented with a tennis ball, for instance, might be said to stand in a certain normative relation with the belief that there is a tennis ball in one's vicinity. The relation is normative in the sense that the latter provides a reason for the former. This, however, is not the sense of intentional states being normative that we are interested in.

The kind of normativity of intentional states that interests us is more basic. We are interested merely in the idea that some intentional states can be evaluated semantically. The notion of semantic evaluation is here to be understood as an umbrella term that covers several criteria of evaluation that correspond to several kinds of intentional states.

Judgments, for instance, are commonly said to aim at truth. To judge that p is to judge that p is true. As such, then, a judgment that p can be semantically evaluated as true or false. If p is true, then one's judgment that p is true, and otherwise false. Desires, on the other hand, are commonly said to aim at satisfaction. To desire that p is to desire that p be satisfied. If p is satisfied, then one's desire is satisfied, and otherwise dissatisfied. Truth and satisfaction, hence, are two instances of the kinds of semantic norms that can be applied to instances of intentionality.

That said, not all instances of intentionality seem to be semantically evaluable. Merely thinking about Obama, for instance, is not semantically evaluable insofar as there is no room for something

to go wrong in one's thinking about Obama. Accordingly, semantic normativity is a central feature of intentionality, common to most, but not all instances of intentionality.

Now, once again, semantic normativity is not a feature of intentionality that can simply be read off the FIS. It is a feature that we discover as we analyze instances of intentionality. It is also a feature that makes intentionality even more mysterious. Indeed, where does this normativity come from? Is it a fundamental feature of intentionality? And what relation does it share with the other features of intentionality? These, again, are questions for a constitutive theory of intentionality.

6.3.4 Intentional Identity

Limit cases set aside, the domain of what we can bear intentional states about is a domain of shared things. In other words, we can generally think about the same things. If John thinks about Obama and Sam thinks about Obama, then both think about the same person. This is the feature we shall call intentional identity.

Identity, however, is a complex word. John and Sam, for instance, drive identical cars. However, we do not mean by this that they share a numerically identical car. We mean, rather, that they each possess a distinct token of a same type of car. In other words, their cars are not numerically but qualitatively identical. What about thoughts? In the case of thoughts, the two notions of identity seem to apply. If both John and Sam think about Obama *as* Obama, then it seem that they both token a thought of a same type: an Obama-thought that is about the same person, namely Obama. Which sense of identity is the one at play in "intentional identity"?

The answer to this question is, as we shall see later, not neutral. What seems to be intuitive enough, however, is that there is one

sense of intentional identity that is non-negotiable. It is the one that we *at least* sometimes think about numerically identical things. If both John and Sam think about Obama, then they are thinking about a numerically identical person. Numerical intentional identity is a non-negotiable feature of intentionality. If we lose it, we face solipsism: a view that no sane philosopher would willingly endorse.

The problem with numerical intentional identity is how far it can be stretched. As we shall see, when combined with some further features of intentionality, the idea that we can think about numerically identical things is an idea that we might have to give up at least in some cases.

6.3.5 Generality

The second variable of the fundamental intentional schema, "*x* is about *y*," can sometimes be filled up by an indefinite noun phrase, as in the sentence, "I desire a beer." To desire something certainly is an instance of intentionality: to desire something is to hold a desire about something. Quite strangely, however, it may be that something one is intentionally directed upon is nothing in particular: one does not desire any particular beer but just a beer, whichever beer it is. This is the penultimate feature we shall be concerned with: generality.

To get the gist of generality, compare instances of the fundamental intentional schema with other schemas, such as, for instance, "*x* shakes hands with *y*." As it appears, it is impossible for this schema to have instances in which one is shaking hands with no one in particular. Of course, one may say, "John is shaking hands with a man," but in this context the indefinite noun phrase, *a man*, works as a referential expression that indicates that the identity of the man with whom John shakes hands is either unknown or irrelevant. This, however, is not equivalent to the claim that John is not shaking

hands with no one in particular. Whenever a hand is shaken, it is a particular hand of a particular person that gets the shake. The intentional schema, hence, seems to possess instances that set it apart from different, though very similar, ones.

This raises some important and puzzling questions about the metaphysical nature of what lies behind the fundamental intentional schema. On the one hand, "x is about y" looks very much like, say, "x is shaking hands with y". On the other hand, there are ways to satisfy the fundamental intentional schema that are ruled out for similar, though non-intentional, schemas. This is likely to leave one wondering about the metaphysics of what lies behind the fundamental intentional schema. How exactly we should then interpret the fundamental intentional schema so that it can account for generality without losing sight of what makes instances of the FIS that instantiate generality intentional is a further desideratum for a constitutive account of intentionality.

6.3.6 Non-Existence

The last major feature of intentionality we shall be concerned with, non-existence, is probably the one that has been the most discussed by theorists of intentionality. Indeed, independently of one's theoretical outlook on intentionality, non-existence constitutes a puzzling phenomenon.

The problem can be formulated as follows. Intentionality as characterized by the fundamental intentional schema naturally seems to be a relation: One thinks about something, one shakes hands with someone—both things seem to be instances of a relational schema. Moreover, it seems more than plausible that a relation, like any predication, either monadic or polyadic, is a determination of what exists. That my desk is white is a determination of my desk, which

exists or has being. Shaking hands, similarly, is a determination of that portion of reality composed of a pair of human beings that exist. And, following Malebranche, where there is no desk or no human being, there cannot be any such determinations: "nothingness has no properties" (Malebranche 1997: 10). Hence if intentionality is a relation, then, *mutatis mutandis*, the same should be also true of intentionality.

At this point, however, things start to go askew. What seems to be true of desks and hands shaking does not mesh well with what we would like to say about intentionality. Indeed, about the claim that intentionality is a relation—understood as a determination of reality or of what exists—Michael Dummett writes the following in the French edition of his book, *The Origins of Analytic Philosophy*:

> This brings up the evident objection that a genuine mental act could exist even without any object in reality: I can fear, or rejoice of, something purely, and, in particular, I can be the victim of a hallucination or any other sensory illusion. All these examples, however, keep the characteristic of intentionality, because a perceptual or auditory illusion does not consist in seeing or hearing without seeing or hearing anything. To possess this characteristic, it is

to be directed upon an object; but, by hypothesis, in that case, there is no object. (Dummett 1991: 47—own translation from the French)[7]

The idea is the following: Intentionality has many features that makes it look like it is a relation. This claim, however, must face the objection that an intentional state that is about some x could exist and remain pretty much the same whether or not x exists. That is, if intentionality is a relation, it seems to be a relation that is a determination of a pair of things of which only one must exist; or, alternatively, a relation that is a determination of an amputated pair of things—*i.e.*, a pair that, paradoxically, is made of only one thing.

Let us take an example. Imagine, for instance, that I am going down the unlit corridor of an old alpine hotel. I distinguish a form in the dark that seems to walk towards me. This causes in me an acute episode of fear. This fear, moreover, is intentional, *i.e.*, it is fear directed towards that form walking towards me. Now, as it happens,

---

[7] The editorial history of Dummett's *The Origins of Analytic Philosophy* is, to put it mildly, rather circumvoluted. The basis of the text was a series of lectures given in English at the University of Bologna. The lectures were then translated into German and published by Suhrkamp under the title *Ursprünge der analytischen Philosophie*. A French translation of the German translation of the English original (sic) was then published in France by Gallimard in 1991. The lectures were then published in their original English, but only after a substantial revision by Dummett, with noticeable alterations to the original text translated and published in German and French. This explains why the above quote is my own translation of the following passage of the French edition: "Cela suscite l'objection évidente qu'un authentique acte mental pourrait exister même sans objet dans la réalité: je peux avoir peur ou me réjouir de quelque chose de purement imaginaire, et surtout je peux être victime d'une hallucination ou de tout autre illusion sensorielle. Cependant ces exemples conservent la caractéristique de l'intentionnalité, car une illusion de la vue ou de l'ouïe ne consiste pas à voir et entendre simplement sans rien voir ou entendre. Posséder cette caractéristique, c'est être orienté vers un objet; mais, par hypothèse, dans ce cas il n'y a pas d'objet." (Dummett 1991: 47)

nothing was walking towards me. What seemed to be there was in fact a mere shadow produced by the light of the moon shining through the window at the end of the corridor. Should we infer from this that I in fact was not the bearer of an intentional state of fear? Or that my intentional state of fear, despite appearances, was not in fact about anything? As Dummett points to, it seems natural, though paradoxical, to claim that in such cases I simply underwent an episode of fear directed upon something that in fact did not exist.

The problem of non-existence, then, seems to face the theory of intentionality with the following problem: If intentionality is indeed a relation, then it seems to be a special kind of relation—one that, contrary to other relations, is not a determination whose relata need all exists. If this is the case, however, one might well wonder, first, how relations can obtain between relata that belong to different ontological realms: existence and non-existence. Second, one might well wonder what it means to say that a relation is, for at least one of its half, is a determination of something that does not exist.

On the other hand, one might reject the claim that intentionality is a relation. This would prevent one from having to settle the weird metaphysical questions just sketched. But what must one then do with the initial pressure to treat intentionality as a relation? If I can fear something that does not exist, this does not prevent my fear from being about that thing. How then can one reject the claim that intentionality is a relation and nonetheless respect these appearances? Either choosing one side of this disjunction or fending it off with a third alternative constitutes another part, if not the main part, of a constitutive theory of intentionality.

7. Conclusion

The overall aim of this chapter was two-fold: first, expound the view that a theory of intentionality should be conceived as a theory of intentional experiences, and, second, introduce our basic assumption about what shape a theory of intentional experiences should take.

With respect to the first aim, we defended the thesis that a theory of intentionality should be a theory of intentional experiences against two main objections: the general objection that the phenomenal cannot be intentional, and, second, the more restricted objection that there cannot be cognitive intentional experiences. As we concluded, these two objections are unwarranted.

With respect to the second aim, we distinguished between two levels at which a theory of intentionality can be pursued: a constitutive and an engineering level. We made clear that the level at which a theory of intentional experiences should be located foremost is the constitutive one. We then introduced the idea that the main aim of such a constitutive theory of intentionality should be to provide an interpretation of what we call the Fundamental Intentional Schema (FIS): a schema that characterizes all instances of intentionality at a maximal level of generality. We finally introduced six main features of intentionality that an adequate interpretation of the FIS should be able to account for.

In the upcoming chapters, we turn to the issue of providing an adequate interpretation of the FIS. The discussion is structured as follows. Chapters 2, 3, and 4 discuss two models that try to implement the idea that the underlying metaphysical nature of intentionality is relational. Chapters 5 and 6 then turn to two further models that deny this claim. The next chapter addresses a first kind of relational theory we call The Intentional Object Model.

Chapter 2: The Intentional Object Model

> Every mental phenomenon is directed toward an object. This thesis sums up such undeniable facts as that there is no idea which is not an idea of something or other, that there is no desire which is not the desire for something or other, that there is no belief which is not the belief in something or other, that there is no seeing which is not the seeing of something or other, and so on. Once pointed out, the thesis of intentionality is a truism. But its philosophical ramifications, as we shall see, are monumental.
>
> (Grossmann 1984: 33)

1. Introduction

In the previous chapter, we introduced two fundamental ideas. First of all, a theory of intentionality should be conceived as a theory of intentional experiences. Secondly, a theory of intentionality conceived as a theory of intentional experiences should essentially amount to an adequate interpretation of what we called the fundamental intentional schema (FIS). This chapter starts down the path of providing such an interpretation.

In terms of their general structure, the upcoming chapters discuss two different families of theories of intentionality: those that argue that the underlying metaphysical nature of intentionality is relational and those that claim otherwise. This chapter focuses on a first relational

theory that we have deemed the intentional object model (IOM). The discussion of the IOM then continues in Chapter 3, which considers some of the IOM's philosophical and logical consequences.

The general idea behind the claim that intentionality's underlying metaphysical nature is relational is that to be the bearer of an intentional state essentially involves standing in a relation to something. As we shall see, however, this idea can be implemented in different ways. In particular, this claim should not be identified with the claim that aboutness itself is a relation. The two claims are conceptually distinct, and the claim that aboutness itself is a relation represents only one means of cashing out the intuitive idea that intentional states are about things.

In this chapter, we concentrate on the main tenets of the IOM. The IOM makes use of two fundamental ideas. First of all, it relies on the idea that aboutness itself is a relation. In other words, the underlying metaphysical nature of intentionality is that of a relation, and this relation is precisely the one of being about something. If I think about Obama, for instance, then I stand in a relation to Obama, the relation of "being about." Secondly, the relata of this relation are intentional objects. Hence, if I stand in "being about" relation to Obama, then Obama is the intentional object of my intentional state of thinking about him. These two claims form the conceptual foundations of a theory of intentionality, according to the IOM. We shall see exactly how and why the IOM articulates them, and we will also explore how they must be completed for the IOM to be able to account for the six main features of intentionality introduced in the second part of the previous chapter.

The chapter is structured as follows: we begin by discussing the idea of a relation itself and then present the main arguments in favor of the claim that the FIS's fundamental nature is relational. Secondly, we detail how the IOM conceives of intentional relations and intentional objects, and formulate a first interpretation of the IOM, the basic IOM

schema. Thirdly, we point to the fact that this schema is too basic to account for the six features of intentionality we identified in the previous chapter and then complement it as necessary. As we shall see, however, some of these features place pressure on the IOM to endorse disputed theoretical commitments. We then turn to this issue in the next chapter, describing these commitments in more details as well as some of their logical implications. That chapter also raises some objections against both these theoretical commitments and the IOM in general.

2. What is a Relation?

What is a relation? Basically, relations are a certain type of properties, namely, polyadic properties. This characterization, however, only brings us to the next question: What are properties? We shall assume that properties are ways of characterizing particulars, that is, features that numerically distinct particulars can share and by which they can be grouped. Sam, for instance, is a human being, and so is Linda. Hence, there is a property that Sam and Linda share and in virtue of which they can be grouped, along with many others, as humans.

We have here a set of notions that are illuminated in terms of each other. A particular is something that can have a property, while a property is an element that characterizes a particular and that forms a criterion for grouping different particulars together. We shall not say much more about either of these notions. In particular, this thesis does not address the perennial question of the nature of properties. Rather, what we require is a metaphysics that allows us to utilize these notions. However, we shall remain agnostic about which particular theory it is.

We shall distinguish between two kinds of properties. On the one hand, monadic properties, such as "being human" in the previous example, play a specific role: characterizing particulars. That is, they allow us to enumerate the characteristics of a specific particular, *e.g.,* Sam, and

permit us to state what properties this particular shares in common with other particulars.

We do not, however, always aim at characterizing and describing single particulars. Indeed, sometimes our objective is to provide characterizations and descriptions focused on how particulars interact with each other. For instance, imagine that John slams his door. One could then describe John as satisfying the monadic property, "slams a door (x)." We might have reason, however, to describe John's slamming his office door not in terms of a monadic property but, rather, in terms of the following dyadic relation: "Slams(x, y)", where John can be substituted for *x* and his office door for *y*.

In that respect, an interesting idea is that not only particulars but also *n*-tuples of particulars can have certain characteristics and, hence, that numerically distinct *n*-tuples of individuals can be grouped in virtue of sharing certain characteristics. For instance, the ordered pairs {John, John's office door} and {Linda, Linda's office door} can both be characterized in terms of the same relation: "Slams (x, y)." Note, however, that the idea here is not merely that *n*-tuples themselves can have properties, which is certainly true but unimportant at this point. It is, rather, the idea that there are genuine and interesting ways to ascribe characteristics to particulars not only on their own but also with respect to the way they interact with each other.

Let us take a further example. Sam is 1.75 m tall, and John is 1.80 m tall. These are characteristics that they possess individually and in virtue of which they can be grouped with other individuals with whom they share the same height. However, taken as the ordered pair <Sam, John>, Sam and John also possess further characteristics that they cannot be taken to have when considered separately. Namely, that pair, along with many other similar pairs, can be characterized as one in which the first member is shorter than the other.

By means of this example, we can then clarify two important ideas mentioned above. First of all, it is individuals, not *n*-tuples of individuals, that instantiate relations. Indeed, it is not the *n*-tuple <Sam, John> itself that is "shorter than." Rather, it is Sam who is shorter than John. However, this characteristic does not hold true for Sam on his own. It is a characteristic of Sam as a member of the *n*-tuple comprised of Sam and John. Moreover, it is not simply Sam and John as a pair that matters for it to be true that Sam is taller than John. It is the ordered pair composed of Sam and John, in that order. In other words, the ordered pair <Sam, John> makes the true the relation "*x* is taller than *y*," while the ordered pair <John, Sam> makes it false. Order, then, matters for relations.

Secondly, the idea of interaction was used above to speak of the relation between John and his office door. The example involving John and Sam, however, allows us to see that "interaction" should not be interpreted too literally when discussing relations. As we use the term, forming an *n*-tuple is sufficient for things to interact with each other. If Sam is shorter than John, this fact alone does not mean that Sam and John interact in any strong sense of the word. In fact, Sam and John might live on different continents, never have met, and even fully ignore each other's existence. Yet, the proposition that Sam is shorter than John would still be true. Why should we call this an interaction, then? We do need not to but the idea is intuitive, and it is hopefully more helpful than damaging.

Particulars and properties can be structured into a hierarchy. At level 0 are particulars. These are not properties but they can instantiate them. At level 1 are properties of particulars, and they are thus called first-order properties. John, a level-0 entity, can then instantiate the property of being tall, which is a level-1 property. Properties themselves, however, can also instantiate properties. The property of being tall, for instance, instantiates the property of being a measurement property.

Being a measurement property, however, is not the kind of property that an individual can instantiate. Thus, such properties must belong to a different level, level-2. Level-2 hence, contains properties of properties, or second-order properties, and so on. This hierarchical conception does not play too much of a role in this chapter, but it does reappear in later chapters and should thus be kept in mind.

Properties of all kinds, however, are generally taken to present the following intuitive, key feature: A predication cannot be true unless that of which it is predicated exists. This trait of properties does play a role in this chapter. As Malebranche states, "nothingness has no properties." (Malebranche 1997: 10) If the red fairy flying around my computer screen does not exist, then she cannot be flying. Applied to the type of properties of which we are currently focusing—relations—this motto yields the following: If at least one of the relata of a relation does not exist, then that relation cannot obtain. If John fails to exist, then Sam cannot be shorter than him. As Uriah Kriegel puts it:

> …just as we do not think that a monadic property can be instantiated in the absence of a particular that instantiates it, so we should not think that a relation could be instantiated in the absence of particulars that instantiate it. (Kriegel 2007: 311)

As we shall later see, this claim, which we shall call the "Malebranche principle," is central to the relational conception of intentionality, as some theories of intentionality, such as the IOM, are bound to deny it. For the time being, however, we shall not enter into a debate about its plausibility, as the next chapter tackles this issue.

Now that we know a bit more about relations, how should we interpret the idea that intentionality is a relation? Of course, we cannot yet express this idea too precisely, as this would force us to make certain theoretical choices we would like to reserve for later. We can, however,

roughly express the gist of the idea as follows: As we saw, intentionality can be characterized as the mind's capacity to be about things. The idea that intentionality is a relation is thus simply the notion that this capacity is relational. In other words, to think about something is to interact with it in the metaphorical sense mentioned above. As we shall see, the entity to which we stand in relation can be that very thing we are intentionally directed upon as in the IOM. On the other hand, it could be something else that we are not intentionally directed upon but that permits us to explain what it is that we are thinking about, as in the content model, which we discuss in Chapter 4.

Taking a bit of distance, we might describe the general outlook of a relational conception of intentionality as follows : Take the case of John thinking about Obama. What shall we say about it? Here is, in very broad terms, how relational theories of intentionality conceive of this fact. First of all, it is true of the plurality of things that constitute what we call reality that two of them, John and Obama, are to be found among them. They also are interacting in a very specific and rather low-key way: The first thinks about the second. In other words, this is not a feature of John taken in isolation. Instead, it is a feature of John as part of an *n*-tuple that also contains, at least, Obama. This seems fair enough. Described in this manner, the existence of intentionality would just be another platitude about reality: there are particulars (first platitude); these particulars have characteristics (second platitude); and particulars interact with one another in an ordered way (third platitude). Intentionality, then, would then just be a particular kind of instance, or set of instances, of the third platitude. This, of course, is a very schematic way of putting things and the question, of course, is how much of this simple idea will survive once we plunge ourselves into the messy reality of intentionality.

3. Relational Intentionality

In the preceding section, we introduced the general concept of relations and quickly sketched one methodological reason for treating intentionality as a relation. Namely, regarding intentionality as a relation would constitute a crucial step towards treating intentionality as a philosophical platitude. However, additional arguments also support the claim that intentionality is a relation. In this section, we briefly present four additional reasons: a grammatical one, a phenomenological one, an epistemic one, and, finally, an inferential one.

3.1 Grammatical Relationality

A first kind of reason that seemingly supports the idea that intentionality does possess an underlying relational structure has two components. First, the FIS unifies all instances of intentionality. Indeed, despite the fact that not all sentences used to describe intentional states share a transitive grammatical structure, a bit of reflection shows that all sentences used to describe intentional states, including those that possess an intransitive grammatical strucure, nonetheless describe states that are instances of the FIS.   Take, for instance, the following two examples:

(1) John thinks about Obama.
(2) John thinks that Obama is American.

Despite of the fact that, grammatically speaking, (1) is a transitive sentence and (2) is an intransitive one, it nonetheless makes sense to claim that both describe intentional states that are instances of the FIS, namely "*x* is about *y*." Indeed, both (1) and (2) are about something, namely, Obama.

The second step is simply to observe that the FIS does possess a transitive grammatical structure that instantiates a dyadic relational schema. From then on, one could conclude that there is reason to infer that the schema's underlying metaphysical nature is also relational. As Reinhardt Grossmann puts it:

> Consider, for example, the sentence 'I see a tree' and compare it with the sentence "A is greater than B." The first sentence seems to mention three things—namely, a person, a tree, and a mental act of seeing. But so does the second; it, too, seems to mention three things, namely A, B, and the *relation* of being greater than. This similarity between sentences about (two-term) relations and sentences about mental acts is rather striking. In fact, it is so striking that it all but forced upon philosophers the view that mental acts are in some sense relational. (Grossmann 1965: 39)

This reason, however, can at best be a *prima facie* one. Indeed, if there is no outright reason to believe that grammatical and metaphysical structures always take separate ways, there also is no general ground to trust that grammatical reasons are conclusive reasons. Philosophers and linguists alike are often prone to remind us that we use expressions such as, "Let's accept this premise for the argument's sake," without there being anything like "the argument's sake". Hence, the claim can at best be that unless one is presented with defeaters that argue against the claim that the grammar and the metaphysics of the FIS go hand in hand, one is justified to believe that this is the case.

3.2 Phenomenological Relationality

The previous reason was purely grammatical. One could complement it with the phenomenological observation that it is not only the FIS's grammar that speaks in favor of the claim that intentionality is a relation.

Additionally, the phenomenology of instances of the FIS also does so. Graham Priest, for instance, claims that:

> … when one fears something, one has a direct phenomenological experience of a relation to the object of the fear. (Priest 2005: 57-58)

This second reason, just like the first one, is a *prima facie* one. That is, there is no reason to definitively doubt the data of experience. Nonetheless, one may be forced, when presented with defeaters, to duck them in favor of other considerations.

3.3 Epistemic Relationality

It seems hard to truly disbelieve that reality is somehow within our cognitive reach. That is, it seems difficult to truly disbelieve that we, at least sometimes, truly perceive, and judge of, mind-independent things that we can reach through intentional states like perceptions and judgments. The question, however, is how can intentional states allow us to thus reach reality? According to some philosophers, the best explanation is that intentionality is a relation. For instance, Reinhardt Grossmann writes:

> To assume that mental acts involve relations to such mind-independent entities as trees and chairs, is to account for the fact that the mind can reach out and make contact with what is not mental. It would account for the apparent fact that we can know not only our own minds but also the world around us. Any other view can either cannot explain this obvious fact or must deny it outright. As has often been noted, *realism*, as opposed to *idealism*, stands or falls with the relational view of mental acts. (Grossmann 1965: 40)

As Grossmann constructs it, his argument for epistemic relationality is directed against idealists or skeptics who claim either that there is no mind-independent reality of which we can have knowledge or, alternatively, that if there is such a reality, it would always remain beyond our reach. According to Grossmann, the realist answer to these challenges is to straightforwardly argue that intentionality is a relation.

Nevertheless, that argument can be, and has been, constructed in different ways A first way—one that seems to be pointed at by Grossmann in the above quote—, would appeal to epistemic relationality in the form of a transcendental argument that would go somehow as follows:

1) The mind can reach out and make contact with what is non-mental.
2) The mind could not reach out and make contact with what is non-mental if intentionality were not a relation.
3) Hence, intentionality is a relation.

Under the assumption that 1) is bedrock, the substantial premise of this argument is, of course, 2). According to 2), it is a conceptually or metaphysically necessary condition for our minds to be able to reach out for the world that we construct intentionality as a relation. The same conclusion, however, can be defended by a weaker kind of argument, namely, an abductive one. Such a line of reasoning would flow as follows:

4) The mind can reach out and make contact with what is non-mental.
5) The best explanation for the fact that the mind can reach out and make contact with what is non-mental is that intentionality is a relation.

6) Hence, intentionality is a relation.

This second version is, of course weaker, than the first. Indeed, the fifth premise does not make it impossible, either conceptually or metaphysically, for the mind to be able to reach out to the world without intentionality being a relation. It merely claims that the likeliest underlying theory of intentionality that would allow us to account for this fact would be a relational one.

This third reason seems much stronger than the first two. Indeed, neither the transcendental argument nor the abductive one delivers mere *prima facie* reasons. They seem to deliver much stronger reasons, as what is at stake in their conclusions is much more substantial than grammar or phenomenology. Indeed, as Grosmann emphasizes, some philosophers have claimed that our ability to mentally reach out to the world is what is at stake in the claim that intentionality is a relation.

3.4 Inferential Relationality

The last reason we should mention appeals to the idea that intentional subjects sometimes perform inferences about intentional experiences by means of using sentences that describe such intentional experiences. More precisely, it appeals to the idea that the best way to give a theory of the validity of these inferences is to cash out the underlying metaphysics of intentional states in relational terms.

Let us, for instance, have a look at the two following sentences:

(1) John thinks about Obama.
(2) Sam thinks about Obama.

These two sentences are descriptions of intentional states, those of John and Sam. Theorists of intentionality commonly posit that these two

sentences can also constitute the premises of a syllogism whose conclusion would be:

(3) Hence, there is something about which both John and Sam think.

Such a syllogism can, of course, also occur in the case of so-called propositional attitudes. For example:

(4) John thinks that Obama is American.
(5) Sam thinks that Obama is American.
(6) Hence, there is something that both John and Sam think.

Theorists of intentionality then face the task of accounting for the intuition that these inferences are instances of a valid schema. It is here that relationality kicks in. Indeed, both (1)–(3) and (7)–(9) seem to be instances of the following schema:

(7) T(a, b)
(8) T (c, b)
(9) ∃(x) T(a, x) & T(c, x)

Granted that first-order classical logic governs the schema (7)–(9) and its instances, the soundness of each instance of this schema requires that the individual constants a, b, and c be well-defined. Only under this condition can then Existential Generalization—namely, the rule of inference that says that from a sentence of the form F(a) one can infer that ∃(x) F(x)—be applied to yield the conclusion.

The claim made by this argument, however, is not that conceiving of intentionality as a relation is the only way to account for the validity of these inferences. Rather, if our aim is to account for the validity of these inferences in a literal way and within the boundaries of classical

logic, then relationality is the appropriate choice. This outcome, however, is not arbitrary. Indeed, it seems that any inferential schema that would account for these inferences' validity in a way that would appeal to a logical form remote from the one proposed in (7)-(9) would result in a theory of inference that would be too complex to account for the fact that intentional subjects sometimes do perform inferences such as (1)-(3) and 7)-9) in a casual and correct way.

This argument is also interesting, because it can be, and has been, used to argue in favor of other theses than the one that intentionality is a relation and thus lies at the center of a theoretical nexus. In particular, this argument has been used to support the existence of propositions (see, *e.g.,* Schiffer 2003, chap. 1). Indeed, if (i) the inference (7)-(9) is valid, (ii) classical logic is assumed, and (iii) intentionality really is a relation, then it can be argued that (iv) propositions must exist to serve as the relata of intransitive intentional verbs.

3.5 Summary

The first two models of intentionality that we are going to examine, the IOM and the CM, are relational models of intentionality. In this section, the goal was to put forward some reasons why theorists of intentionality have supposed that the underlying metaphysics was indeed relational. Besides the methodological claim that a relational conception of intentionality would make of intentionality a kind of philosophical platitude, this section distinguished between four additional reasons of different force and inspiration, namely: a grammatical one, a phenomenological one, an epistemic one, and an inferential one. In the remainder of this chapter, we turn to the first implementation of the idea that intentionality is a relation, namely, the IOM.

4. The Basic Intentional Object Model

If the arguments presented in the preceding section are correct, then intentionality indeed is a relation. This claim, however, is a minimal one that merely tells us about the underlying metaphysical nature of intentionality. In contrast, it does not tell us what kind of relation intentionality is or its relata are. This section explores how the IOM answers these two questions. The first sub-section deals with the nature of the relation, while the second sub-section addressees the nature of the relata.

4.1 Being About

If the arguments of the preceding section are on the right track, then intentionality possesses an underlying relational metaphysical nature. That said, one might wonder about the exact metaphysical shape of that relational nature. This question might appear trivial. Indeed, one may simply infer from the FIS and the data in support of it that the metaphysical structure of this relation is the following: "Being about(x, y)." This, we take, is an adequate rendering of how the IOM conceives of the intentional relation and we shall call it thesis (T1):

> **(T1)** The relation underlying the metaphysical nature of intentionality is the one of 'being about';

However, as this chapter and the next one will make clear, this thesis is very far from being philosophically trivial.

A first important point is that conceiving of the intentional along such lines is not trivial insofar as, as mentioned above, many intentional states are described in terms of sentences that do not instantiate a transitive surface grammar. Take the following examples:

(1) John thinks about Obama.

(5) John fears that Obama will soon be out of office.

According to the IOM, both (1) and (5) are instances of the intentional relation conceived as "being about (x, y)." This might be obvious in the case of (1), while the case of (5) requires some reflection. Indeed, according to the IOM, the fact that (5) does not instantiate a surface grammar close to the one of the relation of (1) does not prevent intentional states like (5) from being instances of the relation "Being about (x, y)." Indeed, according to the IOM, (5) is about something, just like (1) is. Namely, it is about Obama. According to the IOM, this is an obvious phenomenological fact. Intentional states like (5) are about or present something just like intentional states such as (1).

 A second important point is that, thus conceived, the intentional relation is a determinable one. That is, intentional experiences are instances of being about something, but we can nonetheless distinguish between different kinds of intentional experiences, such as: mere thoughts, judgements, desires, hopes, and so forth. The idea is that all of these kinds of experiences are instances of a determinable intentional relation, the one of being about something. As Husserl puts it:

> Intentional experiences have the peculiarity of directing themselves in varying fashion to presented objects, but they do so in an *intentional* sense. An object is "referred to" or "aimed at" in them, and in presentative or judging or other fashion. This means no more than that certain experiences are present, intentional in character and, more specifically, presentatively, judgingly, desiringly or otherwise intentional. (Husserl 2001: 98).

A third important point is that by conceiving of the intentional relation as the "being about" one, the IOM considers that this feature of

intentionality that we earlier called "aboutness" is fundamental. That is, aboutness is directly built into the intentional relation and cannot be derived from other, more fundamental features of intentionality. Again, as Husserl explains:

> If [an intentional] experience is present, then, *eo ipso* and through its own essence (we must insist), the intentional 'relation' to an object is achieved, and an object is 'intentionally present'; these two phrases mean precisely the same. (Husserl 2001: 98)

We understand this short passage from Husserl as follows: Intentional experiences are such that the fact that they have aboutness, conceived in terms of the "being about" relation, is a trait that intentional experiences have in virtue of their essence. That is, there is no need to infer aboutness from a further feature of intentionality. It is there right from the start. If an experience is intentional, then it is *eo ipso* an instance of the "being about" relation.

To sum up this first sub-section, we have claimed that the IOM conceives of the relational structure underlying the FIS as the determinable relation of "being about." According to the IOM, both an analysis of the FIS's surface grammar and phenomenological observations support this claim. Crucially, the IOM considers that aboutness is a built-in, fundamental, feature of intentionality. In the next sub-section, we shall turn to the question of how the IOM conceives of the relata of the relation of "being about."

4.2 Intentional Objects

Our second question concerns the relata of the intentional relation. The IOM makes two central claims in that respect. First, it states that intentional experiences must be irreducibly described in terms of the

concept of an *intentional object* and that instances of this concept serve as the relata of the intentional relation. In other words, for an intentional experience to be an instance of the relation of "being about" is for it to be about an intentional object. The second claim posits that intentional objects are schematic objects, that is, objects of a kind whose nature is exhausted by their satisfying a schema. We shall discuss each of them in this sub-section.

Let us start by trying to grasp the fundamental picture underlying the claim that intentional experiences must be irreducibly described in terms of the notion of an intentional object. Getting this initial idea right will turn out essential to understanding much of what follows, which is to say, both the IOM and the evaluation of its alternatives. In the first chapter, we claimed that intentional experiences are appearances. How, however, should we think of appearances? An appearance, we shall assume, is an appearance *of* something *to* a consciousness. Opening my eyes, I am visually appeared as if I am in my office. Similarly, thinking about Obama, I am cognitively appeared as if I am thinking about Obama. The ways we are appeared to in such experiences convey, in a first approximation at least, what our intentional experiences are *of*.

This first point shall be common ground to all theories of intentional experiences. The question, however, is how one should explain what is meant by the claim that the ways we are appeared to convey what our intentional experiences are of. In other words, what does it mean for an appearance to be *of* something? What does it mean to say that when I think about Obama, I am cognitively appeared as if I am thinking about Obama? According to the IOM, that an appearance is *of* something means that a certain object or topic is presently occupying the focal point of a certain stretch of my conscious life. Whatever that is thus occupying this focal point is called an *intentional object*.

It is essential to clearly understand the strategy of the IOM at this point. Presumably, when I am thinking about Obama, then Obama

himself, the man in the flesh, is what I am thinking about. Accordingly, the IOM is not saying that when I am thinking about, *e.g.*, Obama, I am not related to Obama himself but to a different entity, an intentional object that is numerically distinct from Obama. Rather, the IOM stipulates that we call whatever we are thinking about an intentional object. For instance, if I am thinking about Obama, the man in the flesh, then he is the intentional object of my intentional experience. The notion of an intentional object therefore does not denote a particular kind of thing. Rather, it intends to capture a certain commonality among intentional experiences: They all are about something, *i.e.*, they all have an intentional object. This is the core content of the thesis (T2):

> **(T2)** The relata of the relation of "being about" are intentional experiences (or their bearers) and intentional objects.

Since the word "object" has more than one meaning, it might help to belabor this point a bit further to prevent any misunderstanding. As Tim Crane points out (Crane 2014: 115), a prominent use of the word "object" treats it as synonymous with the words "thing" or "entity," meaning something like "physical particular," *e.g.*, things like trees, tables, chairs, boats, and so on. This, however, is emphatically not the meaning of the word "object" at work in the expression "intentional object." According to the IOM, an intentional object is not—or not only—a physical particular that one is thinking about. Some examples of intentional experiences make this obvious:

> (6) John is impressed by Gödel's first incompleteness theorem.
> (7) John wonders whether justice is an objective value.
> (8) John believes that World War II ended in 1957.

Each of these examples are indisputable examples of intentional experiences. Each of them, however, plausibly fails to be about an object in the sense of a physical particular. According to the IOM, what we shall conclude from this is the negative claim that "object," in the expression "intentional object," is not synonymous with "physical particular."

What, however, is the IOM's positive claim about the meaning of the word "object" as occurring in the expression "intentional object?" One might get to it by pointing to other contexts of use of the word "object," like in the expression "object of attention." As Jerry Valberg remarks (Valberg 1992), in such a context, the word "object" cannot be replaced by "thing" or "physical particular." Indeed, many more things beyond physical particulars can be the objects of our attention. In the expression "object of attention," the word "object" hence seems to serve as shorthand for the longer expression "whatever our attention is focused on." In other words, objects of attention seem to be schematic objects in the sense that an object of attention is whatever can be substituted for the second variable in the schema "x focuses her attention on y."

The role of the word "object" in the expression "intentional object" is very close to its role in the expression "object of attention." Indeed, as Crane notes:

> This is the key to the idea that being an intentional object is not being any kind of thing. For 'intentional object' in this respect (unsurprisingly) is like 'object of attention' rather than 'physical object.' (Crane 2014: 115)

We can express this point in a general way: Like objects of attention, intentional objects are schematic objects in the sense that their conditions of individuation are schematic ones, *i.e.*, for something to be

an intentional object just is for that thing to play a certain substitutive role in an intentional schema, presumably the FIS.

A direct consequence of the fact that intentional objects are schematic objects is that there is no metaphysically substantial kind "intentional object," *i.e.*, objects that have, as constitutive part of their nature, the property of being an intentional object. As Alberto Voltolini puts it:

> Intentional objects are simply objects *qua* objects of thought, *i.e.*, targets of intentional states. Thoughts can be directed at concreta of any kind, as well as at abstracta of any kind – individuals as well as events, numbers as well as fictional entities. This situation requires that intentional objects must be schematic entities, *i.e.*, items which, unlike the above kinds of entities, have no particular intrinsic nature. (Voltolini 2006: 436)

The exact contours of this idea are not easy to grasp, so let us try to unpack it a bit. An object's nature is understood here as a general nature, *i.e.,* what makes it the case that a certain object can be individuated as a token of a certain type *T*. If John is, *e.g.*, a human being, then there is a set of properties that constitutes John's nature *qua* human being. Now, the claim at stake is that whatever nature something that happens to be an intentional object has, it does not at all depend on that thing being an intentional object—that an object is being thought about does not determine its nature. In other words, whatever the nature of something that happens to be an intentional object, this nature is determined completely independently of its status as an intentional object—or, wholly independently of the context of it being the object of an experience. In that sense, an intentional relation leaves the nature of an intentional object untouched.

Consequently, according to the IOM, intentional objects are neither a special kind of object nor is the nature of something affected as a direct result of someone thinking about it[8]. Intentional objects are nothing but objects of intentional experiences, where this is not taken as meaning that they constitute a kind of object but, rather, that they play a role in the individuation of an experience. Metaphorically speaking, an intentional object is everything upon which a conscious gaze, either sensory or cognitive, falls.

If needed, this claim can be further complemented by the observation that the attribution of the property "being an intentional object" to something cannot be the result of a discovery. That is, one cannot inquire about an object and discover that it is an intentional object. That one is thinking about it is sufficient for the property of being an intentional object to belong to that object. Moreover, regardless of whether that property dawns on that object at some point, it remains completely untouched. Its nature is always determined separately from the intentional relation.

Thus far, the dialectic has run as follows: Each intentional experience has an intentional object, and this object can belong to any possible metaphysical category. Hence, an intentional object, like all objects, can have a nature, but this nature is determined outside of the intentional relation. Being an intentional object is therefore not a

---

[8] The distinction between direct and indirect here means the following. Of course, someone can think about something and, as a result, substantially affect it but one cannot do this through the mere power of one's thought. Rather, one needs to think about that thing and act upon it in some way in order to change it. For instance, looking at a stack of wood-planks, one may think about a certain chair to make out of them and might, as a result, build such a chair. That said, one needs to perform a certain number of actions *in addition* to thinking about the chair in order to build it. This also applies to the distinction between thinking something and putting that thought into words so as to produce an effect on someone, *e.g.*, the formation of a belief.

constituent of an object's nature. Objects *qua* intentional objects are merely schematic objects.

However, the issue gains an additional layer of complexity when we consider the fact that attributing a *single* intentional object to intentional experiences seems to be rather arbitrary. Indeed, intentional experiences are complex beasts. For instance, it often happens that we are not merely thinking about, say, Obama but, rather, thinking about Obama in some way. Right now, for instance, I am thinking that Obama will soon be out of office. What, then, is the intentional object of my experience? Is it Obama? Is it the fact that Obama will soon be out of office? Is it the property of being soon out of office? What would constitute a faithful description of how I am appeared to in such a case?

Any restricted answer to that question will seem arbitrary. Assume, for the sake of the argument, that in the case of thinking that Obama will soon be out of office, the intentional object of that experience is the fact that Obama will soon be out of office. Is it not obvious that this experience is *also* about Obama and the property of being out of office? The only way to resolve this problem is to refuse to identify the intentional object of this intentional state with any of the candidates we have mentioned. Rather, the intentional object of an experience is the overall correlate of the experience.

In his paper "Intentional Objects," Crane compares the notion of intentional object with another notion pregnant in the phenomenological tradition, the one of a "world" (Crane 2014: 117). The idea here is that each intentional subject has a world. This claim, however, is not a metaphysical one; it is not a plea for a radical form of relativism, in which each subject lives in its own reality. The idea, rather, is that though intentional subjects target the very same reality, this reality can be disclosed differently to different subjects—or, for that matter, to the same subjects at different times. The idea, then, is not that one experiences a different world, conceived as, for instance, a different

mereological sum of things. Instead, one experiences different features of the same mereological sum, disclosing different worlds at different times. In that sense, an intentional subject *a* pointing out some features of a perceptual scene to an intentional subject *b* is a way for *a* to make her world accessible to *b*. By doing this, however, *a* is not giving the keys of a different world to *b*. She is, rather, making *b* sensitive to certain features of a world we all share. This is the content of the thesis (T3):

> **(T3)** Intentional objects are the overall correlates of intentional experiences, *i.e.*, the intentional object of a particular experience is the world, in the phenomenological sense of that word, as it is disclosed by a particular intentional experience.

An intentional object, in that sense, can be identified with the totality that is disclosed in a particular experience. In the example discussed above, Obama himself, the fact that Obama will soon be out of office, and the property of being out of office are all elements of the world that is disclosed to me as I undergo the experience of thinking that Obama will be out of office. They are the overall correlate of that experience— its intentional object.

The idea, then, is not that we use the concept of an "intentional object" to individuate objects in any metaphysically substantial sense; the individuation conditions of intentional objects, in that sense, are merely schematic, as defined above. The idea, rather, is that the IOM uses it to individuate experiences in a substantive way. Namely, the IOM uses it to claim that there at least some experiences such that we cannot say what these experiences are without mentioning that they have an intentional object, in the stipulated sense indicated in this sub-section.

On the basis of this exposition of theses (T1)-(T3), we shall now formulate a first interpretation of the FIS and put it to the test of the

five main features of intentionality discussed in the previous chapter. We shall then see that an additional thesis needs to be added.

4.3 The Basic IOM Schema

We now wish to bring together the observations from the previous two sub-sections in order to formulate a first interpretation of the FIS, as conceived by the IOM. We call this initial interpretation the basic IOM schema. It is basic, because it stems from preliminary reflections on the intentional relation. As such, it already aims to account for the first of intentionality's main features, aboutness. However, it remains to be seen whether this basic schema can also account for the other primary features of intentionality that we previously distinguished.

The basic schema is the following:

Basic IOM schema: "*x* is about *y*" amounts to "*x* stands in the relation of 'being about' with an intentional object *y*."

Is this schema *too* basic to permit the IOM to account for the other main features of intentionality? The following section addresses that question.

5. Is the Basic IOM Schema too Basic?

In the previous section, we introduced the IOM's two central ideas: (i) the intentional relation is the relation of "being about," and (ii) the relata of this relation are intentional objects. Taken together, these two ideas yielded the basic IOM schema. We now explore the following question: Can the basic IOM schema account for the main features of intentionality that we introduced in the previous chapter? This section addresses each of those features in turn.

5.1 The IOM and Aboutness

The first feature of intentionality requiring an explanation is undoubtedly the most central feature of intentionality: namely, aboutness. As mentioned, aboutness is that feature by virtue of which intentional states can be said to be about things. As we saw above, the IOM treats this feature as a fundamental and primitive one, as it cashes out the underlying relational nature of intentionality in terms of the "being about" relation. An intentional state, then, has aboutness merely by virtue of being an instance of that relation.

This appeal to primitiveness is based on phenomenological grounds. Indeed, according to the IOM, phenomenological experiences harbor a most salient phenomenological trait: They are about something, in the specific sense that they stand in the relation of "being about" with an intentional object. Moreover, this relation is established on the basis of their essence. Husserl, for instance, is particularly explicit on this point:

> We take intentional relation, understood in purely descriptive fashion as an inward particularity of certain experiences, to be the essential feature of 'psychical phenomena' or 'acts', seeing in Brentano's definition of them as 'phenomena intentionally containing objects in themselves' a circumscription of essence, whose 'reality' (in the traditional sense) is of course ensured by examples. Differently put in terms of pure phenomenology: Ideation performed in exemplary cases of such experiences – and so performed as to leave empirical-psychological conception and existential affirmation of being out of account, and to deal only with the real phenomenological content of these experiences – yields us the pure phenomenological generic Idea of *intentional experience* or *act*, and of its various pure species. (Husserl 2001: 96).

According to Husserl, first of all, the intentional relation can be understood in "purely descriptive fashion," that is, as a matter of the way our experiences are given to us. This particularity is not, however, common to all experiences. We have access to this particularity of some experiences, he claims, by means of a "circumscription of essence," that is, by means of seeing what invariantly belongs to an experience by undetarking a series of variation on it. For instance, one may take an intentional experience and make it vary along several lines without losing its intentional nature. According to Husserl, however, what would make it fall outside of the intentional boundaries would be the removal of its instantiation of the intentional relation.

According to Husserl, moreover, it is also a fact accessible through inward examination of our intentional experiences that they must be described by the notion of intentional object as described above:

> Intentional experiences have the peculiarity of directing themselves in varying fashion to presented objects, but they do so in an *intentional* sense. An object is 'referred to' or 'aimed at' in them (*ist in ihnen 'gemeint,' auf ihn ist 'abgezielht'*), and in presentative, or judging, or other fashion. This means no more than that certain experiences are present, intentional in character and, more specifically, presentatively, judgingly, desiringly, or otherwise intentional. (Husserl, *LU*, V, §11)

It is important, however, not to read too much into that claim. From that perspective, the English translation of "*gemeint*" as "referred to" is misleading, as not all instances of something that is "*gemeint*" are instances of reference. Put in our terms, Husserl's claim essentially concerns the meaning of the word "about." As he makes clear, moreover, this "circumscription of essence" is a purely phenomenological matter. That is, it is a claim about the way that experiences are given to us, not a claim about the way that experiences

may, say, relate us to something. As such, then, the claim that aboutness is a relation to an intentional object and that we come to know this fact through phenomenological inquiry is not about reference, understood here as the relation that may obtain between singular terms and existing things. In other words, this claim is not one about the specific mechanism of reference implemented by intentional states. It is, rather, a claim about the meaning of the word "about," understood as characterizing the fundamental feature of intentionality.

That said, however, critics of the IOM might allege that this is precisely one of the features that makes the IOM difficult to accept. Indeed, they might struggle to accept that there is such a fundamental relation of being about that cannot be explained in more fundamental terms. In other words, detractors might not concede that reality has an irreducible feature, the nature of which is to be about other things, in a way that is not explainable in other terms. We next turn to two lines of answer against this critique.

First of all, it is not clear why a descriptive account of aboutness cannot be explanatory. Of course, it is may not be explanatory in the same sense than, say, reduction to more fundamental concepts is, but it might nonetheless be an acceptable explanation of aboutness to be told that it is a relational feature that is invariably present in any intentional experience and that it constitutes its essence.

Secondly, one might also appeal to the principle of neutrality between the constitutive and the engineering task introduced in the previous chapter. The constitutive level has been described as a theory of intentionality *as such*. That is, it primarily focuses on the different concepts needed to describe the phenomenon of intentionality and the various relations between them. The IOM's claim regarding aboutness is a *constitutive* claim in the strict sense. It might have some incidence on the engineering level, but this is not yet an argument against it *per se*. Of course, if two accounts are completely equivalent at the constitutive level

but one of them simplifies the engineering task, then this counts as an argument against the account stuck at the constitutive level. However, as mentioned, this requires two accounts that are equivalent at the constitutive level, and the IOM's analysis of aboutness indicates that this is precisely what is not available. According to the IOM, a correct description of intentional experiences is one in terms of the notion of intentional object, and of intentional experiences being relations to them.

5.2 The IOM and Aspectuality

The second feature requiring an explanation is aspectuality. By this, we mean that intentional states always contain a certain embedded point of view. Consequently, what an intentional state is about is never targeted from, to use Nagel's famous expression, "the view from nowhere." (Nagel 1986) Rather, such a state is always targeted from a certain point of view, the one occupied by the intentional subject. As we saw, moreover, this notion of point of view can either be taken somewhat literally, as in the case of perception, or more metaphorically, as in the case of cognitive acts.

In Chapter 1, we used examples to demonstrate the existence of aspectuality as a feature of intentional states. Let us examine one more as a reminder. We start with the following two instances of intentionality:

(9) John is thinking about Phosphorus.
(10) John is thinking about Hesperus.

Aspectuality is that feature by virtue of which (9) and (10) are said to be different intentional states. Indeed, even though (9) and (10) are intentional states that are about the same thing, since "Phosphorus" and "Hesperus" are co-referential terms, it is nonetheless the case that (9)

and (10) are different intentional states. This is made clear by the possibility for John to regard (9), but not (10), as a correct description of his mental life. That is, if John ignores that "Phosphorus" and "Hesperus" are co-referential, then he may well deny thinking about Hesperus while assenting to thinking about Phosphorus.

Aspectuality poses a problem for the IOM. Indeed, we claimed earlier that the IOM asserts that intentional objects are merely what an intentional subject is thinking about, as determined from the outside of the intentional relation. From this, it follows that what John is thinking about is the same thing in (9) and (10), namely, that thing to which both "Phosphorus" and "Hesperus" refer. In other words, taken as instances of the basic IOM schema, there simply cannot be a difference between (9) and (10), because the terms of the intentional relation are exactly the same in both cases. Hence, this schema, which says that intentionality is an unmediated relation between an intentional experience—or its bearer—and an intentional object, cannot account for aspectuality.

According to the IOM, aspectuality demonstrates that there is no such thing as an intentional state that is about an intentional object as such, that is, an intentional object in its entirety, exhaustively, and with all of its aspects present at the same time. On the contrary, an intentional object is always intended under a certain aspect. Venus, for instance, can be intended either *as* Phosphorus or *as* Hesperus, but it cannot be intended as such. As Tim Crane puts it:

> So, at its most general, the idea of aspectual shape is just the idea that there is no such thing as a thought about, or an awareness of, an object *as such*—that there is no such thing as what we might call 'bare' presentation of an object. (Crane 2001: 20)

The challenge for the IOM is thus to modify the basic IOM schema to make it strong enough to be able to account for this idea. The question is

what needs to be added to that first interpretation of the FIS to permit the IOM to account for aspectuality. Basically, what seems to be needed is a way to implement a fundamental distinction between two things, namely: (i) an object as it is, or as such; and (ii) an object as it is intended or presented. Indeed, according to the IOM, it is beyond question that John is thinking about the same thing in both (9) and (10), namely, that thing to which "Phosphorus" and "Hesperus" both refer. Nevertheless, it also seems unmistakably clear that this very same thing is intended differently in (9) and (10). Subsequently, what is needed is a means of accounting for this distinction.

At this point, one might wonder whether the notion of intentional objects itself really is insufficient to account for the idea of an "object as intended." That is, one might speculate as to whether one could simply answer that John is thinking about different intentional objects in (9) and (10). Note, first of all, that this would not lead to a means of accounting for the above distinction between (i) and (ii); instead, it would conflate the two. Secondly, it should be obvious that such a way out would be incompatible with the above characterization of intentional objects. Indeed, as we defined it, intentional objects are merely what they are from outside of the intentional relation. In other words, this solution would simply not work, because it claims that the intentional relation gives rise to a new kind of object, namely, objects as intended. This is

incompatible with the IOM's understanding that intentional object are determined from outside of the intentional relation[9].

According to the IOM, the mistake committed by this proposal is that it believes that a new kind of object, distinct from our common-sense worldly objects, objects as intended, must be introduced to deal with this feature. According to the IOM, there is another, less problematic way to account for aspectuality. This approach is to distinguish, for each intentional state, between two different things: (i) its intentional object, namely, what it is about; and (ii) its intentional content, namely, the particular way or aspect under which it intends its intentional object.

In the case of the example just discussed, the IOM then claims that we must differentiate between two things. On the one hand, there is that thing to which both "Phosphorus" and "Hesperus" refer and that would constitute the intentional object of the intentional experience in both cases. On the other hand, there is the particular way John is thinking about this thing in (9) and (10), that is, the intentional content of these states. The intentional object is what transcends the intentional state and what it points to, while the intentional content is, in metaphorical terms, the specific configuration of the mind as it intends an intentional object or the specific way a certain object appears,

---

[9] Note that some accounts of intentionality that we will not discuss in these pages have proposed to treat intentional objects in just that way. Franz Brentano's account of intentionality in his *Psychology from an Empirical Standpoint*, for instance, is sometimes interpreted in that way (cf., *inter alia*, Chrudzimski 2014). Such accounts, however, should be avoided because, first, they do not respect the phenomenological fact that what are intentionally directed upon are objects themselves, and not objects specially designed to be the relata of the intentional relation. And, second, such accounts must be able to give us a theory of the relation between the intended objects, as we may call them, and transcendent, real-world, objects. This, however, constitutes an unnecessary difficulty.

sensorily or conceptually, in an experience. In the case of (9) and (10), John can intend the same object while his mind is in two different configurations, the "Phosphorus" one in (9) and the "Hesperus" one in (10).

The idea of a "configuration of the mind" is highly metaphoric but it is, hopefully, intuitive. The general idea is that there might be different ways to do the same thing, *i.e.*, different ways to think about the same thing. This idea should of course be made much more precise. In particular, much should be said about its diversity. Indeed, it is likely that different kinds of intentional states, such as purely sensory and purely cognitive ones, have different kinds of contents. It is also likely that the an intentional subject's degree of sophistication—*i.e.*, whether it is a human or non-human intentional subject, whether it is an infant or an adult, and so on—do play a key role in the individuation conditions of these configurations. What is important, however, is to point to the fact that, according to the IOM, the introduction of intentional contents does not amount to the introduction of a new kind of objects, in the sense of a new kind of things that intentional states are about. Indeed, intentional contents are not what intentional states are about. Rather, they are about intentional objects, but there are different ways of being about the same intentional object.

This introduces an important modification to our above discussion of the conditions of individuation of intentional experiences. We said that intentional experiences are individuated in terms of what they are about, *i.e.*, their intentional object. However, the existence of aspectuality demonstrates that this cannot be the full story. Indeed, according to the IOM, two distinct experiences—thinking about Phosphorus and thinking about Hesperus—will be individuated as tokens of different kinds of experiences. But this could not be the case if intentional experiences were individuated solely in terms of their intentional objects, as these two experiences are about the very same

thing, namely, that to which "Phosphorus" and "Hesperus" co-refer. What the introduction of the notion of intentional content reveals, according to the IOM, is that intentional experiences must be individuated more finely than by their intentional object. They must be conjointly individuated by their intentional object and their individual contents. Thinking about Phosphorus and thinking about Hesperus are tokens of dissimilar kinds of experiences. Despite the fact that they are about the same thing, they are tokens of different types of intentional contents—a Phosphorus content and a Hesperus content.

We now clarify four additional points about the notion of content. While some of these are very general and should apply to any notion of content, others concern the IOM's specific means of implementing that notion.

First of all, the general idea of intentional content is not proprietary to the IOM. The way that the IOM conceives of intentional content is certainly original, but the idea that something like a notion of content should be introduced for a theory of intentionality to be capable of accounting for aspectuality is not unique to the IOM. Each notion of content, however, must face the same worry. In a nutshell, intentional objects and intentional contents are conceptually distinct notions, but we do not possess expressive means for properly distinguishing between them. This may give rise to some difficulties in getting a firm grip on the distinction.

For instance, when asked about the intentional *object* of (9), *Phosphorus* is a correct answer. But *Phosphorus* would also be a correct answer if one inquired about its intentional *content*. This, of course, does not mean that we have provided the same answer in both cases. An intentional content is not, apart from exceptional cases in which one is thinking about intentional contents, an intentional object. At this point, philosophy hits the barriers of language. Even though there is a conceptual distinction between intentional objects and intentional

contents, this distinction is clothed in an air of ineffability. Horgan, Graham, and Tienson summarize this point as follows:

> The object of an intentional attitude qua intentional object cannot be identified independently of the content. It can only be identified from the inside, as it were, the content. (Horgan, Graham & Tienson 2009: 521)

This point is not new. Commenting on Frege's distinction between sense and reference, which is very close to that between intentional object and intentional content, Michael Dummett underscores the difficulty of stating what the sense of a proper name would be:

> It has become a standard complaint that Frege talks a great deal about the senses of expressions, but nowhere gives an account of what constitutes such a sense. This complaint is partly unfair: for Frege the sense of an expression is the manner in which we determine its reference, and he tells us a great deal about the kind of reference possessed by expressions of different types, thereby specifying the form that the senses of such expressions must take. It is true enough, however, that he says practically nothing about what the senses of different expressions consist in; and this is a legitimate ground of complaint. *Indeed, even when Frege is purporting to give the sense of a word or symbol, what he actually* states *is what its reference is; and, for anyone who has not clearly grasped the distinction between sense and reference, this fact makes his hold on the notion of sense precarious.* The sense of an expression is the mode of presentation of the referent: in saying what the referent is, we have to choose a particular way of saying this, a particular means of determining something as the referent. In a case in which we are concerned to convey, or stipulate, the sense of an expression, we shall choose that means of stating what the reference is which displays the sense: we might

borrow here a famous pair of terms from the *Tractatus*, and say that, for
Frege, we *say* what the referent is, and thereby *show* what its sense is.
(Dummett 1973: 227 – italics added)

The same remark applies to the distinction between intentional objects
and intentional contents. One cannot state what the intentional object of
a certain intentional act is *except in a certain way*. There is no neutral or
absolute way to state what is the intentional object of a certain
intentional experience. Hence, one must always pick a particular means
of identifying the intentional object. By stating what the object is, one
thereby *shows* under what intentional mode this intentional object is
thought about.

The second point concerns the specific way that the IOM cashes
out the notion of content. That is, according to the IOM, an intentional
content is the way in which a certain intentional object is presented.
Hence, the notion of intentional content is defined in terms of the
notion of an intentional object. The idea, however, is not that the former
concept can replace the latter concept. Instead, the notion of intentional
content complements the notion of intentional objects. In other words,
according to the IOM, to fully individuate an intentional state, one must
not only identify what the object is about, *i.e.*, its intentional object, but
also state its intentional content, *i.e.*, the way one intends a certain
intentional object. As Tim Crane puts it:

> I need this distinction [between intentional object and intentional
> content] because neither directedness nor aspectual shape on their own is
> enough to characterize what I mean by the subject's perspective on the
> world. Directedness on an object alone is not enough because there are
> many ways a mind can be directed on the same intentional object. And
> aspectual shape alone cannot define intentionality, since an aspect is by

definition the aspect under which an intentional object (the object of thought) is presented. (Crane 2001: 29)

Formally speaking, the distinction betwee intentional objectsand intentional contents is then treated as the addition of a term to the intentional relation, which becomes a triadic relation. In other words, intentional experiences are no longer conceived as dyadic relations that obtain between intentional experiences, or their bearers, and intentional objects. Rather, they are triadic relations that obtain between intentional subjects, intentional contents, and intentional objects. An intentional object, hence, is intended by an intentional subject under a certain aspect, as captured by the notion of content.

Thirdly, a central distinction between intentional objects and intentional contents concerns the way that we access them. In other words, what is appeared to us in an intentional experience, what the experience is about—its direct object, to use a grammatical label—is its intentional object. In short, awareness of intentional objects is a direct kind of awareness. It is a matter of being aware *of* something. Intentional contents, on the other hand, are not routinely presented to the intentional subject in such a way. That is, according to the IOM, when one thinks about an intentional object, this object appears to the subject. However, intentional experiences do not feature additional appearances of intentional contents.

Of course, one could be presented with an intentional content. This, however, is an ambiguous claim. On the one hand, one could be presented with an intentional content in the sense of being intentionally directed upon that content, *i.e.*, if one turns an intentional content into an intentional object. On the other hand, however, one must recognize a second sense in which one could be presented with an intentional content. That is, by being about an intentional object, an intentional experience is an appearance, sensory or cognitive, of an intentional

object. Admittedly, however, an intentional experience also constitutes an implicit or indirect appearance of an intentional content.

That is, since one cannot be intentionally directed upon something but under some aspect, then for each intentional experience, an intentional object is presented under some aspect. The intentional object, however, is not presented as distinct from the aspect under which it is presented. From the point of view of the intentional subject, there is only one thing one is presented with: the intentional object as it is presented. That is, one is presented with an intentional content, but one is *not* intentionally directed onto that content. In other words, one is intentionally directed through that content but is not intentionally directed upon it. By being intentionally directed through a content, hence, one is *indirectly* presented with it. As John Tienson puts it:

> There can, of course, be presentations of content, but such presentations can only be indirect, as the content through which such and such object is presented. (Tienson 2013: 486)

Explicit or direct awareness of an intentional content can then occur as one realizes that the structure of intentional states distinguishes between intentional objects and intentional contents. At that point, one can hypostatize intentional contents and turn them into intentional objects.

Fourthly and finally, one may choose to discharge this notion of "configuration of the mind," or content, in different ways. However, at minimum, what accounts for this idea of configuration should make room for the notion that contents are non-transparent. By transparent contents we mean the idea that if a same intentional object O can be presented by two different contents C and C', then it is impossible for an intentional subject to ignore that C and C' actually present the very same intentional object. Transparent contents, in other words, wear their intentional object on their sleeve. But if intentional contents were indeed

transparent, we could not use them to account for cases like the that of (9 and (10). In other words, the notion of content must allow an intentional subject to ignore that two distinct contents present the very same object.

In summary, introducing the notion of intentional content permits the IOM to offer a second interpretation of the FIS. This interpretation accounts for both aboutness and aspectuality. This second, more complex interpretation is as follows:

> **IOM schema:** "*x* is about *y*" amounts to "*x* stands in the relation of 'being about' to an intentional object *y,* and for some content *C*, *C* presents *y* to *x*."

Having introduced this second interpretation, the next-subsection evaluates how it fares with respect to the remaining of the features of intentionality we introduced in Chapter 1.

5.3 The IOM and Semantic Normativity

Semantic normativity is the feature of intentionality according to which at least some intentional states can be evaluated as meeting, or failing to meet, a certain semantic norm. The notion of content again plays an important role in explaining the existence of semantic normativity.

Indeed, each intentional object that is being thought about is being thought about in a certain way, as captured by the notion of content. There is, however, no reason to think that contents always match intentional objects' features. That is, there is no justification for believing each way of thinking about an object indicates how that object really is. In that sense, some contents can be, for example, false or inaccurate. According to the IOM, this explains semantic normativity's existence.

Take the following case. Walking by night on the streets of Fribourg, John sees at a certain distance something he takes to be a fox. He then forms the belief that there is a fox in front of him. At it happens, however, what John takes to be a fox is, in fact, nothing but a piece of trash lying on the sidewalk. The content of his intentional state thus presented its intentional object in an inaccurate manner. Namely, the piece of trash in front of him is not a fox, even though it was presented to John as such. The idea, then, is that the IOM possesses the resources to account for semantic normativity in terms of the notion of content.

5.4 The IOM and Intentional Identity

Intentional identity is that feature by virtue of which we can say that we are thinking about a shared world. That is, it permits us to say that we are thinking about the same thing. According to the IOM, this feature is, at least at first glance, quite easy to grasp. Indeed, two intentional subjects can be said to be thinking about the same thing only if they are thinking about the same intentional object. Moreover, granted that the nature of intentional objects is determined from the outside of the intentional relation, then nobody's thinking impinges on what anyone thinks about. There is then no obstacles to the claim that we can think about the same things. As we shall see in the next chapter, moreover, a strength of the IOM is that it can extends this treatment of intentionality to cases in which non-existence kicks in.

5.5 The IOM and Generality

We can think about many different kinds of things. In other words, we can have intentional objects of many kinds. A problem, however, is that

there seems to be more kinds of intentional objects than kinds of things. Generality is an instance of this problem. Let us take the following example:

(11) John desires a beer.

According to the IOM, this instance of intentionality amounts to the claim that a relation of "being about," under the determining guise of desire, obtains between one of John's experiences (or John himself), an intentional object, and an intentional content that presents the object in some way. The problem at hand, however, concerns the nature of the intentional object. Indeed, what kind of intentional object is "a beer?" And how can it satisfy John's desire?

The problem can be expressed more precisely as following. Following the letter of the IOM, it seems that (3) should be analyzed as following: "Desires (John, a beer, [a beer][10])," where "a beer" is treated as an individual constant referring to a particular element of the domain. This, however, raises two problems. First, it does not seem that there is anything in the domain that corresponds to "a beer." That is, there certainly are many particular beers that exist but it does not seem as if a general, or indefinite, beer also exists in addition to these. Particular instances of beer are are all the beers that there are. Second, it seems mistaken to interpret John as needing a particular indeterminate thing to satisfy the desire as expressed in (1). That is, only a determinate beer would satisfy John's desire. It is just that there is no particular beer that is required for his satisfaction, or, as we may put it, the identity conditions of a particular beer do not figure in the conditions of satisfaction of his desire. A such, it seems that applying naively the IOM schema to instances of generality creates troubles for the IOM.

---

[10] Square contents denote the intentional content of the state.

Can the IOM account for generality in some other way? As Chapter 6's discussion of a non-relational version of a theory that irreducibly describe intentional experiences in terms of the notion of intentional object will make clear, there might be a way to simply accept at face value the following conjunction: (i) the intentional object of (11) is literally a beer and (ii) there is no such thing as an indeterminate beer. The relational version of the IOM, however, must find another answer. This solution should not result in the postulation that indeterminate objects serve as the semantic value of indefinite noun phrases, such as "a beer." As mentioned above, this would simply fail to deliver the right conditions of satisfaction of John's desire.

Thus, the only way out for the IOM is to look for a relatum that is determinate and nonetheless possesses the necessary properties required to yield suitable conditions of satisfaction for an instance of intentionality such as John's desire in (11). According to the IOM, this can be done by endorsing the claim that (11) is a description of John's intentional experience that is elliptic for the one that John's desire is, in fact, about the existential fact that he has beer. Accordingly, (11) would then amount to the following:

(12) John desires that he has a beer.

According to the IOM, this would allow us to claim both that John's desire has a relata of a determinate nature, namely an existential fact and not an indeterminate beer, and that John's desire would be satisfied just in case he is handed any particular beer and not an indeterminate object. This is achieved in virtue of the fact that the fact that John has a beer has a the following logical form: $\exists(x)$ Beer (x) & Has (John, x). That fact would then obtain just in case its existentially quantified part gets a witness – that is, any witness that meets the requirements of being a beer.

At this point, one might fear the IOM getting out of its track and committing itself to a conceptual analysis with problematic phenomenological consequences. However, some phenomenological observations support this claim about the logical form of an intentional experience expressed by means of indefinite expressions. Imagine Henry sees a specific ice cream—let us call it Creamy—in someone's hand. Creamy looks delicious, and Henry forms the desire to eat it. That is, he forms the desire to eat the very ice cream he sees in someone's else hands, Creamy, and not anything else. Compare this desire with another one: namely, Henry's desire to eat any ice cream, although formed upon seeing Creamy in someone else's hand. This latter desire is not one for Creamy itself but, rather, for anything relevantly similar to it. In other words, Henry does not want to grab Creamy out of someone else's hand and eat it but, rather, wants to acquire a numerically distinct ice-cream for himself.

Now, it is natural enough to describe the first of these desires as being about a particular ice cream: namely, Creamy. But what about the second? Does it seem natural to describe it as a desire for a particular, yet indefinite, thing? We could plausibly deny this claim by saying that desires expressed with indefinites are primarily about oneself, that is, about facts obtaining about oneself, rather than about indeterminate things. For instance, if Henry desires to have ice cream, then phenomenologically speaking it seems much more plausible that he desires that the fact that he has ice cream obtains rather than that he desires a particular, though indefinite, ice cream.

If these remarks are correct, then generality does not constitutes a problem for the IOM. It is not, however, completely unproblematic. Indeed, unless one believes that all facts, including future ones, exist at all times, understanding Henry's desire for ice cream as a desire for the existential fact of having ice cream seems to be comprehending it as a desire for a fact that does not (yet) obtain. Accordingly, this analysis of

indefinites delivers the IOM both the relationality and the semantic indefiniteness for which it longs. However, by way of compensation, it commits the IOM to the claim that there are non-obtaining facts to which we can stand in relation. In other words, according to the IOM, the problem of indefinites collapses into the problem of non-existence. This is the issue that the next sub-section will address.

5.6 The IOM and the Problem of Non-Existence

The last feature that we shall discuss is non-existence, namely, the fact that intentional states can be about things that exist as much as they can be about things that do not exist. Consequently, by virtue of non-existence, it follows that for anything that we can think about, that thing either does or does not exist. This is a mere fact about intentionality.

The IOM, however, cashes out this fact in a particular way. Indeed, when conjoined with the general application of the IOM schema, the claim that we can think about things that do not exist yields the thesis that some intentional objects do not exist. Accordingly, the IOM does not simply claim that we can think about things that do not exist. It also states that there are things that do not exist, and that we can stand in the "being about" *relation* to them.

This constitutes a difficult challenge. Indeed, it forces the IOM to give us an acceptable account of the literal truth of the claim that "there are things that do not exist, and that we can stand in the 'being about' *relation* to them." This challenge, however, constitutes the main point of disagreement between the IOM and its adversaries. These latter claim that this requires the IOM to endorse a Meinongian theory of objects that, to use Karel Lambert's phrase, cannot "keep the line between fantasy and philosophy."(Lambert 1983: 1)

In the next chapter, we discuss in some details Meinongianism and develop some main lines of criticism against it. We shall then have to suspend our evaluation of the IOM up to this point.

6. Conclusion

In the second chapter, we covered the following ground: First, we discussed some reasons for treating intentionality as a relation. Second, we offered a first, basic, interpretation of the FIS according to which for an intentional experience to be about something essentially is for it to stand in the determinable "being about" relation to an intentional object. Third, we inquired as to whether this first interpretation of the FIS was strong enough to account for the six features of intentionality identified in Chapter 1. We saw that if it could account for aboutness to some degree, it could not account for aspectuality. We then proposed a second, slightly more complex, interpretation of the FIS, which we called the IOM schema. This second interpretation differed from the first in that it introduced a third relata—intentional contents—to the intentional relation. Fourth, we inquired whether this second interpretation was strong enough to account for intentionality's other features. The results, however, were mixed. Indeed, as we saw, this second interpretation stumbles on the problem of non-existence. Indeed, since the IOM schema asserts that all instances of intentionality are relations to intentional objects, it follows from non-existence that some instances of intentionality are relations to non-existent intentional objects. This, however, is a substantial claim and a final evaluation of the IOM must be postponed until after our investigation of the Meinongian claims that: (i) there are objects that do not exist, and (ii) we can stand in relation to them. The next chapter takes up these questions.

Chapter 3: Meinongianism

> A relational account of the intentionality of mental acts, however, runs immediately into the problem of nonexistent objects. This problem arises because the nature of the relation clashes with the nature of intentionality. Relations seem to require existing terms, while mental acts can intend nonexistent objects and circumstances. Can these two features be reconciled? I do not think so.
>
> (Grossmann 1984: 45)

1. Introduction

Some intentional experiences are about things that do not exist. This simply is a fact about intentionality. However, coupled with the claim that all instances of intentionality are instances of the IOM schema introduced in the previous chapter, this fact yields the following substantial conclusion: Some intentional states are relations to non-existent intentional objects. This conclusion commits the IOM to the truth of the following two theses:

(i)   Some intentional objects do not exist; and:
(ii) Non-existent intentional objects can be the relata of the relation of "being about."

Let us assume that the IOM schema correctly captures the nature of intentionality and that the following is true: "John thinks about Pegasus." From this, (i) can be inferred as follows:

1) John thinks about Pegasus;
2) Pegasus does not exist;
3) Hence, some intentional object about which John is thinking does not exist.

The second thesis, on the other hand, simply follows from the claim that intentionality is a relation. That is, if (a) intentionality is a relation, (b) the statement "John thinks about Pegasus" is true, and (c) Pegasus does not exist, then (d) it follows that there is a relation that obtains between John and Pegasus, despite the fact that Pegasus does not exist. This, however, contradicts the intuitive principle that only things that exist can have properties and stand in relation to each other—what we, in Chapter 2, referred to as "Malebranche principle".

Being committed to (i) and (ii), the IOM now faces the challenge of explaining why they are true. In other words, the IOM must be able to philosophically defend both (i) and (ii). In this chapter, we shall discuss the standard means of achieving this, namely, the endorsement of Meinongianism, a body of doctrine initiated by the Austrian philosopher Alexius Meinong and later developed by a series of followers including, Richard Routley (Routley 1980), Terence Parsons (Parsons 1980), and, more recently, Graham Priest (Priest 2005) and Francesco Berto (Berto 2013).

We shall not, however, aim at faithfully interpreting the historical Meinong himself. Nor shall we attempt to discuss any possible version of Meinongianism—a task for an inquiry of its own. Rather, we shall present a general overview of the main tenets of Meinongianism, which will allow us to cover if not all, at least most, versions of it.

The chapter is divided into two parts: an expository and a critical one. The expository part is composed of four sections, each introducing one main tenet of Meinongianism. We first clarify the notion of nonexistence as understood by Meinongians before turning to three central Meinongian metaphysical principles: (a) the principle of indifference, (b) the principle of independence, and (c) the characterization principle. In the critical part, we discuss two kinds of objections, inconclusive and conclusive ones. While the inconclusive objections are not sufficient to dismiss Meinongianism, they nonetheless allow us to better grasp its nature and commitments. From the conclusive objections, however, we can conclude that the endorsement of Meinongianism has problematic consequences that one should rather avoid. Moreover, since the IOM is bound to endorse Meinongianism to be able to account for the above theses (i) and (ii), this creates insuperable difficulties for its interpretation of the FIS.

2. Non-Existence

In this section, we introduce Meinongianism in general terms by means of four theses, (T4)-(T7). We start with the general thesis of Meinongianism, which can be expressed as follows:

> **(T4)** Not all objects about which we can know truths and be intentionally directed upon exist, and some of them do not exist *at all*.

What kind of theory is Meinongianism? That is, what kind of thesis is (T4) expressing? This is a complex question, and Meinongians themselves are not united in their answers to it. Does Meinongianism amount to an ontological theory, or is it something else altogether? That is, is Meinongianism encouraging us to substantially augment our

ontology? Or, in contrast, is Meinongianism, or at least some versions of it, compatible with a rather frugal ontology?

If the first is the case, then Meinongianism can be said to defend the claim that we are ontologically committed to more things that those that exist. In other words, when ontologists compile the complete list of the kinds of things that there are, they must take into account not only those things we know to exist, like mountains, tables, and chairs, but also those kinds of things that do not exist, of which, for instance, the green dragon hovering around the room which I am currently hallucinating is an instance.

If the second is the case, however, then Meinongianism is not a theory of interest to practitioners of ontology—or at least not for the exact same reasons. That is, when Meinongians say that there are things like the green dragon I am now hallucinating, they are not saying that expressions like "there is" or "there are," when utilized to express claims like, "there is a green dragon hovering around the room right now," are used in the same sense as when used to express claims like, "there are mountains in Switzerland."

Our take on these two possible options should be rather opinionated. There is, in fact, a fundamental problem with interpreting Meinongianism as an ontological theory. Ontology is virtually universally defined as the theory of what has being, and Meinongianism—at least in its original form—recommends countenancing not only what exists, in any possible sense of the word, but also what does not exist, in the strong sense of not possessing any kind of positive ontological status *at all*—neither existence, being, nor anything else. Hence, interpreting *de facto* Meinongianism as an ontological theory seems uncharitable. If this is correct, then Meinong is not making an ontological claim when summarizing his view in a famous passage of "The Theory of Objects," saying:

> Those who like paradoxical modes of expressions could very well say: "There are objects of which it is true that there are no such objects." (Meinong 1960: 83)

Rather, under this interpretation, Meinong is pointing to the fact that the two occurrences of the quantifier "There are" are not used in the same sense. The first one is used in a non-ontological way, while the second is used in an ontological way.

Nevertheless, if Meinongianism is not an ontological theory, what else could it be? Our interpretative starting point is that Meinongianism should be interpreted as a theory of quantification. In particular, Meinongianism amounts to a theory of absolutely unrestricted quantifiers and their domain. From the Meinongian perspective, and contrary to a common doctrine discussed below, quantification does *not* constitute the premises of an ontological theory, *i.e.*, a theory about what possess any kind of positive ontological status[11].

First, what is the difference between a restricted and an unrestricted quantifier? Quantifiers range over a certain domain. Moreover, this domain can either be the overall domain of quantification or a contextually restricted portion of it. Restricted quantifiers are thus those quantifiers that we use to range over a contextually restricted portion of the overall domain, while unrestricted quantifiers range over absolutely everything.

Take the following example of restricted quantification: Within the context of a party, someone says, "All the beer has been drunk." This person's statement is true. It is, however, compatible with the fact

---

[11] We shall use the notion of a positive ontological status in order to remain neutral with respect to the distinction between being and existence, endorsed by some Meinongian, such as Meinong himself, but rejected by others, such as Routley (1980) and Priest (2005), as this distinction shall not play any important role for us.

that there is more beer available to purchase at the supermarket. That speaker, however, is not contradicting herself by both saying that all the beer has been drunk and that not all the beer has been drunk. Rather, her first use of the quantifier is contextually restricted to that sub-part of the overall domain of quantification which is the party. When restricted to that portion of the domain, her assertion is correct, but when that contextual restriction is lifted, it turns out to be false.

What about unrestricted quantifiers? Another interpretation of the sentence, "All the beer has been drunk," is available. Namely, someone could utter it with the aim of asserting that not only has the party run out of beer but also the entire universe has gone dry, down to the very last drop. Thus, unlike restricted quantifiers, unrestricted quantifiers are not contextually restricted to a certain portion of the whole domain of quantification. Absolutely everything falls within their range, *e.g.*, every single drop of beer.

According to the interpretation we will defend here, Meinongianism is to be understood as a theory of absolutely unrestricted quantification. It is a theory about what such quantifiers can quantify over, and consequently, it indicates what sentences can be regarded as true. In particular, Meinongianism countenances the claim that the boundaries of our overall domain of quantification are not constituted by the boundaries of our ontology. There is, on the one hand, that to which we are ontologically committed. On the other hand, there is that over which we can quantify[12].

---

[12] That we are not interested in interpreting the historical Meinong should be obvious here as Meinong does not frame his theory within a formal logic apparatus. However, what we mean here is that this is what Meinong does mean once we do express his view from within such a framework. Or, better expressed, what Meinong does mean about quantification in a vernacular language and its best translation into a formal language.

What is the underlying motivation behind Meinongianism? The answer to this question is multipronged, but a key impetus behind it stems from the fact that quantificational expressions *in the vernacular* clearly and abundantly quantify over both existent and non-existent things. Among many possible examples, the following two natural language sentences certainly are grammatically well-formed, commonly regarded as true, and at least intuitively, quantify over things that do not exist (although, as we shall see in §3.7.2, anti-Meinongians have contested this intuition):

(1) Some things do not exist.
(2) Not everything exists.

On top of that, there also seems to be a range of predicative operations on domains that do not seem to presuppose existence, such as:

(3) Pegasus does not exist.

and, granted that Pegasus does not exist:

(4) Pegasus has wings.

Some interpretations of the nature of quantifiers, moreover, treat these as mere devices of generalization. Hence, if (3) and (4) are true, we can infer that some things do not exist and that some non-existent things have wings.

Of course, all of these claims are open to refutation. Nevertheless, according to Meinongianism, there is no way but to take them as our starting point. Indeed, according to Meinongianism, we must take these facts as a sign that the overall domain over which we can quantify is not equivalent to the domain of things that possess a positive

ontological status. Whatever the most generous ontology countenances, it would still be incomplete in terms of the things over which we can quantify and predicate. Quantification and ontology drift apart and, hence, the former is by no means a safe guide to the latter.

Such an account is far from trivial. It flatly contradicts Quine's famous, and now standard, claim made in the opening of his landmark paper, "On What There Is":

> A curious thing about the ontological problem is its simplicity. It can be put in three Anglo-Saxon monosyllables: "What is there?" It can be answered, moreover, in a word—"Everything"—and everyone will accept this answer as true. (Quine 1953: 1)

Indeed, according to Meinongianism, it is *not* the case that everything is there if, like Quine, we understand the question "What is there?," as an ontological one. Indeed, not everything exists and, hence, it is not the case that everything is there in Quine's sense. Indeed, according to Meinongianism, the "overall inventory of what there is," to use an expression coined by Alberto Voltolini (Voltolini 2006: ix)—the domain of absolutely unrestricted quantification—contains both things that exist and things that do not exist. Existence, therefore, does not constitute the outside border of the things over which we can quantify, as Quine has it. Rather, it is a genuine property that can discriminate between things we quantify over. In other words, according to Meinongianism, when we say "everything" in a purely unrestricted way we are not, *pace* Quine, giving an answer to an ontological question but, rather, pointing to everything that can be quantified over, either existing or not.

But what exactly does our overall domain of quantification contain if not the full extent of the ontologist's list? Is there a systematic way to determine what it comprises? According to Meinongianism, there is an answer to these questions, but it is not one which we can reach

through any standard kind of inquiry, *i.e.*, either an empirical or a metaphysical investigation. On the contrary, to answer that question we need a completely different kind of theory which Meinong himself calls a "theory of objects." As he puts it in his most famous essay:

> If we remember how metaphysics has always been conceived as including in its subject matter the farthest and the nearest, the greatest and the smallest alike, we may be surprised to be told that metaphysics cannot take on such a task [of being a theory of objects]. It may sound strange to hear that metaphysics is not universal enough for a science of Objects… For the intentions of metaphysics have been universal (a fact which has often been disastrous to its success). Without doubt, metaphysics has to do with everything that exists. However, the totality of what exists, included what existed and will exist, is infinitely small in comparison with the totality of Objects of knowledge. This fact easily goes unnoticed, probably because the lively interest in reality which is part of our nature tends to favor that exaggeration which finds the non-real a mere nothing – or, more precisely, which finds the non-real to be something for which science has no application or at least no application of any worth. (Meinong 1960: 79)

In this passage, Meinong is very explicit. There is, according to him, something that deserves to be called a "theory of objects" and whose domain is broader than metaphysics itself, which is only concerned with what exists. The totality of what exists, however, is *infinitely* small in comparison to the totality with which the theory of objects is concerned. Indeed, the theory of objects addresses a significantly larger domain of Meinongian unrestricted quantification. As Meinong puts it, this domain contains everything that can be an *object of knowledge*. The following thesis encapsulates this idea:

(T5) The nature of our overall domain of quantification must be determined by a separate kind of science, a theory of objects, whose concern is not everything that exists but, rather, everything that can be an object of knowledge.

"Knowledge" is a rather big word and upon getting acquainted with (T5), one might justifiably take it to be pretty crazy. However, one should not be fooled by Meinong's usage of the expression "object of knowledge." Indeed, he does not mean that one can gain substantial knowledge from a theory of objects—that is, knowledge of the type one could gain regarding things that possess any sort of positive ontological status. Rather, the kind of knowledge relevant to (T5) is, as he puts it, "utterly lacking in natural interest" (Meinong 1960: 82). To put it in other words, what he means by "object of knowledge" is loosely equivalent to "something that can be thought about." However, he eschews the purely psychological talk and is adamant to couch his theory in epistemic terms because he takes it that thinking is, in some admittedly strange way, a way to acquire knowledge of what is in our domain of quantification *anyway*. His use of an epistemically charged language is explained by the fact that to know something about objects in a way that is utterly lacking in natural interest is, according to Meinong, still knowledge of *mind-independent* and *public* objects.

Indeed, according to Meinong, what is to be found in the most general domain of quantification does not depend on us in any sense. On the contrary, the content of this most general domain is determined completely independent of us, and it is up to the theory of objects to *discover* what exactly this most general domain contains. That said, it is true that the starting point of historical Meinongianism lies in the theory of intentionality, as Meinong subscribes to the IOM and takes it for granted that there must be some objects that do not exist. Indeed, as he noted at the beginning of 'The Theory of Objects':

> That knowing is impossible without something being known, and more generally, that judgments and ideas or presentations (*Vorstellungen*) are *impossible without being judgments about and presentations of something*, is revealed to be self-evident by a quite elementary examination of these experiences. (Meinong 1960: 76–emphasis added)

It would be a mistake, however, to interpret Meinong as inferring his distinctively Meinongian theory from his theory of intentionality. There might be, of course, genetic reasons related to intentionality that led him to develop his theory of objects, but it would be a mistake to interpret it as dependent upon his theory of intentionality. On the contrary, according to Meinong, that non-existent objects play a role in the theory of intentionality is just a particular instance of a much more general fact that goes beyond intentionality. This is the fact that, completely independently of us, some things do not exist.

In other words, Meinongianism is neither an ontological theory nor a theory of *intentionalia*, *i.e.,* of objects that would have to be added to the overall inventory of what there is merely to serve as the relata of a sub-class of intentional experiences: namely, those that are about things that do not exist. From a systematic point of view, it is instead the other way around. It is because a proper understanding of the domain of unrestricted quantification allows for non-existent objects in a completely mind-independent way that there can be non-existent intentional objects. As Meinong claims a few pages later in "The Theory of Objects":

> If we were now to maintain the aforementioned subjectivity of sense-qualities, we could speak of the object of a presentation of blue only in the sense of something which is a capacity of that presentation, from which reality withholds, as it were, the opportunity for its realization…

However, *I cannot conceal from myself at present the fact that it is no more necessary to an Object that it be presented in order not to exist than it is for it to exist.* (Meinong 1960: 83–emphasis added)

Hence, a theory of objects is concerned with what Meinong calls "the object taken as such and in general" (Meinong 1960: 78), and taken as such, an object is not mind-dependent. The object taken as such is as independent of the mind as any other commonsensical object. Also, the object as such is not an object understood in an ontological sense, *i.e.*, an object to which we are ontologically committed. Rather, it is an object that is there independently of us, upon which we can be intentionally directed, about which we can know things in Meinong's non-substantial use of that verb, and, finally, of which we can assert, or deny, a positive ontological status. Accordingly, Meinong maintains that such an object is *not* yet an of concern for ontology. This leads us to the following important thesis, to which we shall return in the next section:

> **(T6)** The target domain of a theory of objects are mind-independent—and, especially, intentionality-independent—objects about which we can think and of which we can gain knowledge prior to the determination of their ontological status.

Accordingly, contrary to a widespread misunderstanding, Meinong's account does not collapse into that put forward by the young Russell, who believed for awhile that everything over which we can quantify and about which we can think and talk, existent or not, must have some form of being. As Russell asserts in the *Principles of Mathematics:*

> Whatever may be an object of thought, or may occur in any true or false proposition, or can be counted as *one*, I call a *term*. This, then, is the widest word in the philosophical vocabulary. I shall use as synonymous

with it the words unit, individual and entity. The first two emphasize the fact that every term is *one*, while the third is derived from the fact that every term has being, *i.e. is* in some sense. A man, a moment, a number, a class, a relation, a chimaera, or anything else that can be mentioned, is sure to be a term; and to deny that such and such a thing is a term must always be false. (Russell 2010: 44-45)

In this passage, Russell claims that any constituent of a proposition must possess some kind of being. In contrast, Meinong strictly rejects this doctrine. As he puts it:

Instead of deriving the being of an Object from the being of an Objective… it would be better to conclude from the facts with which we are concerned that…the being of an Objective is not by any means universally dependent upon the being of its Object. (Meinong 1960: 85)

Meinong's notion of an "objective" corresponds to the Russellian notion of a proposition, *i.e.*, objectives are what can be asserted. However, Meinong is very clear: He thinks that an objective, such as *Pegasus does not exist,* can be a true and existent objective, even if all the constituents of this objective do not exist. Indeed, he thinks that the being of an objective is *not* universally dependent upon the being of the object with which an objective is concerned. Hence, Meinong's doctrine is different from Russell's. Not everything over which we can quantify and about which we can talk must have some king of being.

Finally, we claimed above that, according to Meinong, the task of a theory of objects is to *discover* the elements of its domain. Moreover, we also stated that the methodology of a theory of objects was neither the one of empirical sciences nor the one of metaphysics. How, then, should Meinongianism proceed? Meinong's answer is straightforward: we should employ an axiomatic method. That is, we should work our way

towards some self-evident axioms and derive valid conclusions from them—among which, Meinong was adamant, will follow the sound basis for the truth of Meinongianism. In other words, according to Meinongianism, a theory of objects can deliver a systematic answer regarding the nature of our overall domain of quantification. This constitutes the final preliminary thesis:

> **(T7)** The theory of objects axiomatically determines the nature of our overall domain of quantification.

In the next section, we turn to the three main axioms of a Meinongian theory of objects: the principle of indifference, the principle of independence, and the principle of characterization. In the second part of the chapter, we will evaluate the consistency of these axioms with other independently plausible metaphysical and logical principles.

3. The Principle of Indifference

The principle of indifference is arguably the most fundamental axiom governing the theory of objects. It pertains to the nature of objects. We shall not try to give a sophisticated definition of what Meinong means by nature. Instead, we shall merely adopt the following superficial characterization: A thing's nature is that which makes a thing what it is. Any property that makes something that which it is constitutes part of its nature. By ostension: Possessing the property of a human being makes me what I am and is hence part of my nature. By comparison, being currently seated in the library does not make me what I am and hence is not part of my nature.

The principle of independence can be taken as an elucidation of the above (T6):

(T6) The target domain of a theory of objects are mind-independent and intentionality independent entities as they are accessible to us prior to the determination of their ontological status.

Indeed, what does it mean that there are objects accessible to us prior to the determination of their ontological status? Meinong cashes out this idea as follows: Each object has a nature. This nature, moreover, is always determined *Außersein*—literally, beyond being. In other words, we have a strictly *a priori*, rational, access to the nature of things, *prior to* and *independently of* the determination of their ontological status. By virtue of this fact, we can be intentionally directed upon each of them and know truths about them. The most general domain of quantification can hence be identified with the totality of objects whose nature is determined *Außersein*.

To have a nature determined *Außersein*, however, is not to have a peculiar, deflated kind of being. Rather, it is a condition under which we can determine what our overall domain of quantification contains. Note, however, that by determining what there is *Außersein*, we do not gain substantial knowledge about the world, understood as the sum of what does possess a positive ontological status. We only acquire non-substantial knowledge about the content of our overall domain of quantification. Hence, a theory of objects is not an *a priori* inquiry into reality, conceived as the province of ontology. It is something distinct from, and even prior, to that.

Indeed, if the principle of indifference is true, then it is therefore false that, as Russell claims in the above quote, whatever is spoken or thought about needs to have being. There is no need for anything to have any kind of being to be thought about and be an object of knowledge. In fact, from a Meinongian perspective, it is actually the converse that is true: Any statement that pertains to the ontological

status of an object requires that this object has been thought about prior to the determination of its ontological status. As Meinong elucidates:

> Blue or any object whatsoever is somehow given *prior to our determination of its being or non-being, in a way that does not carry any prejudice to its non-being.* We could also describe the situation from its psychological side in this way: if I should be able to judge that a certain object is not, then I appear to have had to grasp the Object in some way beforehand, in order to say anything about its non-being, or more precisely, in order to affirm or to deny the ascription of non-being to the Object. (Meinong 1960: 84–emphasis added)

Meinong's principle of indifference is admittedly radical. As we shall see, however, it is not unprecedented. Presenting one of its ancestors will provide it with some legitimacy and also help us to clarify what is unique to Meinong's understanding of the doctrine of indifference. Thus, we shall first approach this principle through a brief excursus *via* one of its illustrious predecessors, Kant's doctrine of natures as exposed in his famous discussion of the ontological argument in the *Critique of Pure Reason*. After having visited these Kantian origins, we shall return to Meinongianism and to its radical understanding of the doctrine of indifference.

3.1 The Kantian Origins of the Principle of Indifference

As is it well-known, the ontological argument tries to define God into existence. As the most perfect being, he must possess all perfections, *i.e.*, all positive properties. Subsequently, he must exist. The key premise of this argument is that existence is a first-order property like any other. This claim, however, has been long contested by philosophers and, especially, by Kant.

The point is not that existence does not work, grammatically speaking, like a first-order property. Indeed, from a grammatical point of view, "Obama speaks" and "Obama exists" seem to be of the same breed: They are subject-predicate sentences. The point, rather, is that the similarity between existence and other properties does not extend beyond surface grammar. In other words, the claim is that even though existence might have the grammatical appearance of a perfectly standard first-order predicate, this fact does not yet prove that it is one.

The 17ᵗʰ century French philosopher Pierre Gassendi formulated a first possible objection to the claim that existence is a standard first-order property. He argues that existence is not a property but rather the *pre-condition* for having properties (see Persnyk 1998, chap. 2). Accordingly, an ontological argument that treats existence as a property would be mistaken on the grounds of a blunt error of category. The problem is not merely that existence is treated as a standard first-order property but, rather, that it is treated as a property *at all*.

Kant offers a second possible reaction to the claim that existence is a property. In the *Critique of Pure Reason*, he famously claims that the ontological argument is a failure because whatever existence really is, it is not, in any case, a *real* property (Kant 2007). Hence, God cannot be defined into existence in terms of possessing *by definition* the property of existence.

Kant characterizes the notion of a real property as that which determines or characterizes an individual. In other words, according to Kant, a property is real if its instantiation by an object tells us something about the kind of thing that it is. For instance, the instantiation of the property "being square" determines a particular and tells us something about the kind of thing it is: a square thing. Hence, "being square" is a real property. Kant, however, denies that existence is such a real property, *i.e.*, he refutes that existence tells us something about the determinations of an object. It might work grammatically as a property,

but it does not play the role of determining an object, and so it is not a real property. According to Kant:

> "Being" is obviously not a real predicate; that is, it is not a concept of something which could be added to the concept of a thing. It is merely the positing of a thing, or of certain determinations, in themselves[13]. Logically, it is merely the copula of a judgment. The proposition, "God is omnipotent", contains two concepts, each of which has its object— God and omnipotence. The small word 'is' adds no new predicate, but only serves to posit the predicate *in its relation* to the subject. If, now, we take the subject (God), with all its predicates (among which is omnipotence), and say, "God is", or, "There is a God", we attach no new predicate to the concept of God, but only posit the subject in itself with all its predicates, and indeed posit as being an object in relation to my concept. (Kant B627|A599)

Let us break this passage into two parts. In the first part, up to: '…God and omnipotence,' Kant is making a point about the logical structure of predicative propositions. As he sees the matter, they are comprised of three things: two concepts, God and omnipotence, related by a copula, "is." The role of the copula is only to attribute the second concept to the first. Furthermore, it does not attribute a second property to God, namely, existence. Accordingly, if we were to suppress the second concept, leaving us with, "God is," we would not be left with a

---

[13] Here, we modify the Kemp Smith translation to follow the one of Markus Weigelt which, we contend, follows more closely the German original: "Es ist bloß die Position eines Dinges, oder gewisser Bestimmungen an sich selbst." Indeed, to translate this sentence as Kemp Smith does as: "It is merely the positing of a thing, or of certain determinations, as existing in themselves," is to add the unjustified presupposition that to posit a thing, or a certain determination, is to pose it as existing.

predication of existence but with a subject, "God," a copula, "is," and a missing second concept. In other words, "God is," according to Kant, is an incomplete predicative proposition which awaits of a concept to which to relate the concept of God.

In the second part of the passage, Kant then uses this observation to support the claim that whatever the predicate of existence is doing, it cannot attribute a real, determining, property. Rather, saying "God is" or "God exists" amounts to taking the concept of God as it has been antecedently determined and, as Kant puts it, "posit[ing]" it, with all of its predicates. Thus, if existence might be said to function grammatically as a predicate, it does not constitute a real predicate.

Part of the difficulty in this passage is interpreting exactly what Kant means by the act of positing a subject or positing an object in relation to a concept. One approach is to read Kant as asserting that existence claims are not about particulars but about concepts. In other words, asserting that God exists does not amount to asserting of God himself that he exists but, rather, that the individual concept of God has an instance.

One way to support this interpretation is by supplementing this first passage with a second one, located just a few paragraphs down:

> Through whatever and through however many predicates I may think a thing (even in thoroughly determining it), nothing is added to it if I go on to say that the thing is. Otherwise, what exists would not be the same as was thought in the concept, but would be more; and I could not say that the exact object of my concept exists. Even if I were to think in a thing all reality, except one, that one missing reality would not be supplied by my saying that so defective a thing exists, but it would exist with the same defect with which I thought it – for otherwise what exists would be different from what I thought. (Kant B628, 629|A600, 601)

In this passage, Kant seems to be saying that we must distinguish between two things: the determination of the nature of a thing and the determination of its ontological status. The first is a matter of thinking what is "in the concept," while the second amounts to asserting that some existent object corresponds to that concept. However, existence cannot be found among those properties that determine the nature of something. Indeed, if this was the case, then we could not say that an antecedently determined nature has an instance, since asserting its existence would simply further determine its nature.

Of course, this interpretation is controversial, but it does not seem to completely lack textual evidence (for a different—and much more thorough and expert—interpretation of Kant's view on existence, see Rosefeldt 2011). It can be summarized as follows:

1) A property is defined as that which characterizes or determines an object;
2) We can think, *i.e.,* have general or individual concepts, of things that do not exist;
3) If it is assumed that existence is a determining predicate, then any concept of a thing that does not exist cannot have an instance, because an attribution of existence would not result in the positing of an antecedently determined object but in an act of further determination;
4) Hence, existence is a not a real predicate.

Ultimately, this may or may not be Kant's position[14]. It is one, however, with textual evidence. Moreover, it is one that allows us to draw an

---

[14] At least some of Kant's interpreters understood him that way, such as Jonathan Bennett, who writes: "According to Kant, every existence statement says about a concept that it is instantiated, rather than saying about an object that it exists." (Bennett 1974: 231)

important contrast with Meinong. Also, even if its conclusion is merely negative, it is nonetheless strong enough not only to reject the ontological argument but also Gassendi's view, as mentioned above. That is, if this argument is sound, then it is possible for there to be fully determined natures of things that do not exist. Kant puts forward his own view of the matter in terms of concepts: There can be a fully determined concept of something without there being any instance of that concept. Pegasus, for instance, can have a fully determined nature, in the sense of there being a fully determined individual concept of Pegasus, independently of any consideration of his ontological status.

Hence, Kant seems to be adhering to some form of a principle of indifference, *i.e.*, there can be a thing's nature, understood as a fully determined concept of that thing, independently of the existence of something with that nature. There can be a concept of Pegasus—or, for that matter, of God—independently of the attribution of any kind of positive ontological status to that which it is a concept of. In that sense, Kant agrees with Meinong that the nature of a thing is indifferent to its existence. Meinong's own principle of indifference, however, is much more radical than the one we just attributed to Kant. The next sub-section makes this clear.

3.2 Meinong's Radical Understanding of the Doctrine of Indifference

Meinong's view is conceptually related to that of Kant, but it is distinctively stronger. The following passage summarizes Meinong's own view about nature's indifference to being:

> …neither being nor non-being can belong essentially to *the object in itself*… This may also be expressed in the following less engaging and also less pretentious way, which is in my opinion, however, a more

appropriate one: *the object* is by nature independent to being. (Meinong 1960: 86–emphasis added)

Meinong's view on the principle of indifference amounts thus to the following: by nature, objects are independent to being and non-being. We can think about them and know their nature independently of their possessing any kind of positive ontological status, as their nature is determined *Außersein*. As such, an object that is the subject matter of the Meinongian theory of objects is an instance of what Routley calls a "pure object" (Routley 1980: 857).

A first striking difference between Kant and Meinong with respect to the doctrine of indifference concerns the level at which Meinong locates his principle of indifference. According to Meinong, it is objects that may not exist and have nature and not, as in Kant's view, concepts that may not have instances. This is a major difference that underlines the distinction between a theory of concepts *à la* Kant and a theory of objects *à la* Meinong. According to Meinong, the principle of indifference does not only mean that there can be concepts of things that do not exist. Rather, there can be objects just like any others—that is, objects with natures, including objects of thought and knowledge— that simply do not exist. For Meinong, this follows from the fact that for any object, the determination of its nature occurs *Außersein*, that is, prior to and independently of the determination of its ontological status.

Among the objects whose natures are determined *Außersein*, Meinong thinks that we can make a three-fold distinction. First of all, there are objects that actually have a positive ontological status, *i.e.,* that have either the property of existence or the property of subsistence. Secondly, there are objects that neither exist nor subsist but *could* do so. And, finally, there are objects that can *neither* exist nor subsist. Indeed, according to Meinong, an "absurd object," as he calls it (Meinong 1960: 86), such as Mill's round square or Twardowski's steel cannon made of

wood, can be said to have a nature that is determined *Außersein* even though it "carries in itself the guarantee of its own non-being in every sense" (Meinong 1960: 86).

This last sentence permits us to make the following important remark on the principle of indifference and its relation to ontology. Meinong, of course, does not deny that some objects do possess a positive ontological status. Moreover, his doctrine of indifference is even compatible with the claim that some objects *must* exist or *cannot* exist. Indeed, as Findlay (Findlay 1963) points out, it is essential to understand the doctrine that objects have their nature determined *Außersein* to see that in all instances Meinong distinguishes between two things: the nature of an object and what *follows* from this nature. Accordingly, what the principle of indifference claims is that existence or non-existence never figures into a thing's nature. However, nothing prevents a thing's nature from determining whether it can, cannot, or must exist. As such, some objects can have their nature determined *Außersein* and nonetheless be necessary, contingent, or impossible objects. However, from the perspective of the theory of objects, which considers all objects *Außersein*, all such objects which fall into different metaphysical categories all stand on par as pure objects. Any statement that pertains to the ontological status of an object is a metaphysical or ontological venture beyond the boundaries of the theory of objects.

This points to a second important difference between Kant and Meinong. Indeed, Kant claims that a thing's nature can be fully determined independently of its existence. Nonetheless, he does not thereby endorse a Meinongian doctrine according to which the nature of a pure object is determined *Außersein*. Indeed, according to Kant, concepts do possess a positive ontological status, being, which falls short of existence (*cf.,* Perszyk 1993, Chap. 2). Hence, according to Kant, if it is true that natures are indifferent to the existence of what they are natures of, this does not yet mean that they are completely indifferent to

any kind of positive ontological determination. Meinong's doctrine, on the other hand, is that all objects can have their nature determined *Außersein*, independently of any kind of ontological determination at all. This is most clear in Meinong's recognition that not only possible objects but also impossible ones, such as the round square, have their nature determined *Außersein*. Indeed, in Meinong's eyes, the round square is genuinely round and square, even though, as we saw above, its contradictory nature makes it necessarily non-existent.

Finally, a third possible source of difference concerns the status of existence. Indeed, if there is some plausibility in the claim that Kant interprets existence as a second-order predication, Meinong himself definitely interprets it as a first-order one. Consequently, from a Meinongian perspective, even if we cannot conceive of existence as a property that determines a thing's nature, we can nonetheless conceive of existence as a "discriminative property." (Rosefeldt 2011: 333). Indeed, it allows us to distinguish, among those things over which we can quantify, between those that exist and those that do not.

To summarize the present discussion, Meinong's principle of indifference amounts to the claim that all objects have their nature determined *Außersein*, that is, independently of their having any kind of positive ontological status. Accordingly, an object can have a nature even if does not possess any kind of positive ontological status at all. Its nature, in other words, is indifferent to being. We saw, moreover, that this principle might be said to have historical roots in authors like Kant, even though Meinong's interpretation of the doctrine of indifference turns out to be much stronger. In the next section, we focus on a second principle, the principle of independence. In Meinong's words, this principle "presents a welcome supplement to the principle that the subject is by nature indifferent to being" (Meinong 1960: 86). To what extent? Let us now turn to that question.

4. The Principle of Independence

4.1 The Question of Predication

A first confrontation with the principle of indifference might be unsettling. One might wonder, for instance, in virtue of what can an object have its nature determined *Außersein*. Indeed, a plausible intuitive principle ruling the relations between properties and particulars—what we earlier called the Malebranche principle—is that an object cannot instantiate a property without possessing *some* kind of positive ontological status. It is at this point that Meinong's principle of independence kicks in. Its job, in a nutshell, is precisely to give support to the notion that predication does not entail existence, and hence, that non-existent objects can have their nature determined *Außersein*.

Meinong's own formulation of the principle of independence is as follows:

> The *Sosein* of an object is not affected by its *Nichtsein*. The fact is sufficiently important to be formulated as the principle of independence of *Sosein* from *Sein*. The area of applicability of this principle is best illustrated by consideration of the following circumstance: the principle applies, not only to objects which do not exist but also to objects which could not exist because they are impossible. Not only is the much heralded gold mountain made of gold, but the round square is as surely round as it is square. (Meinong 1960: 82)

As a first pass, this second principle can be formulated as follows: An object's instantiation of properties is independent of its having any kind of positive ontological status. The golden mountain, for instance, does not exist. Still, according to Meinong, this fact does not prevent the golden mountain from instantiating both the property of being golden and the property of being a mountain. In Meinong's words, the *Sosein* of

a pure object, *i.e.*, its being some way or another—is independent of its *Sein*, *i.e.*, its having a positive ontological status.

Bearing in mind the above intuitive principle that predication entails existence, one might immediately wonder how can Meinong claim that properties can be instantiated by things that do not exist. This, however, is probably the wrong dialectical point for gaining a proper understanding of Meinong's doctrine of independence. Indeed, from a Meinongian point of view, it would be mistaken to begin from a certain conception of properties and then to move on to the claim that properties cannot be truly predicated of non-existent objects. Indeed, according to Meinongianism, we must take our cue from a more neutral standpoint and mould our notion of what properties are in such a way that the principle of independence holds. This more neutral standpoint, according to Meinong, is the *semantic* operation of predication itself.

Indeed, one could distinguish between two levels at which a theory of properties can be discussed: a metaphysical one and a semantic one. At the metaphysical level, a theory of properties is concerned with the conditions under which a particular can instantiate a property. At that level, it primarily focuses on answering three questions: (a) What is a property?; (b) What is a particular?; and (c) What is instantiation, the relation that obtains between a property and a particular?

At the semantic level, on the other hand, a theory of properties is concerned with a related but slightly different question: Under what conditions can a predicative sentence or proposition be true? At this semantic level, a theory of properties turns into a theory of *predication,* and the above questions (a)-(c) fade into the background to be replaced by queries such as: (d) Which sentences can be regarded as predicative?; and (e) Under what conditions can such sentences be evaluated as true?

Now, having distinguished between (a)-(c) on the one hand, and (d)-(e) on the other, one could either start with answering the former and use one's results to answer the latter. Or, alternatively, one could start

with the latter and use one's results to answer the former. Meinong, under our interpretation, fully embraces the latter option.

In other words, Meinong's motivation for endorsing the principle of indifference is not located at the metaphysical level. Rather, he is inspired by semantic concerns. The semantic side of the principle of independence amounts to the claim that almost all sentences that look like predicative sentences can and should be treated as such, independently of the ontological status of whatever subject matter of which they appear to be predicating something. Much to the dismay of anti-Meinongians, the Meinongian school of thought can thus be interpreted as taking this fact as a *requirement* on the metaphysical level of a theory of properties.

At the metaphysical level, Meinong then takes himself to be forced to extend a traditional realist picture of properties to account for the fact that non-existent objects can instantiate properties as much as existent ones. Accordingly, the principle of independence claims that instantiation is a *cross-ontological* relation—it can obtain between entities of different ontological categories, such as between properties that presumably exist and particulars that do, or do not, exist. Moreover, these properties that can be instantiated by non-existent objects are perfectly normal properties, *i.e.*, the very same properties that can be instantiated by both existent and non-existent objects. For instance, the non-existent golden mountain and the existent golden papal ring share an identical property, being made of gold, as the most recent quotation testified. The next section will then be devoted to discussing why and how Meinong imposes this order of things.

4.2 The Traditional Theory of Predication

Let us begin with the following question: What is a predicative sentence? Thus far, we have spoken about such topics without clarifying whether

we meant sentences in a natural language or a formal language. Presumably, however, we could ask the question of what is a predicative sentence in both cases. That is, we could ask which natural language sentences are predicative, or alternately, which natural language sentences can be interpreted as predicative in a formal language of choice.

The first option, however, could be interpreted as meaning one of two things. First of all, we could interpret it as an inquiry about the *grammatical* form of predication in natural language sentences. Secondly, we could understand it as an inquiry into the *meaning* of natural language sentences that feature a predicative grammatical form. The first kind of analysis is not philosophically crucial. Indeed, in purely grammatical terms, we know what predication is in natural languages. It is the joining of a name with an adjective by means of a copula, such as in "John is happy." What we would like to know, however, is what such natural language sentences actually *mean*, *i.e.*, what are the propositions that such sentences express.

An examination of the propositions expressed by natural language sentences requires us to interpret natural language sentences into some formal language. Thus, understood in a philosophically substantial way, the first of the above two options collapses into the second one. Knowing which natural language sentences are predicative is knowing which of them express propositions that are predicative. The issue, then, is not the grammatical form of predication *per se*. Instead, it is the *logical* form of a predicative sentence and the relation between the former and the latter. We shall then focus on this latter issue.

There is, however, no general answer to the question, "What is the logical form of a predicative sentence?" Indeed, the answer will always be restricted to a target formal system L or a family of such systems. Thus, the appropriate way of phrasing this question is, "What is the logical form of a predicative sentence in a formal language L?" Different formal languages are likely to give divergent answers.

We shall begin by discussing a first theory of predication, what we call the traditional theory of predication (TTP). We shall mention how the TTP intends to carve out the boundaries of the notion of a predicative sentence and discuss how this notion of predicative sentence puts pressure on the interpretation of natural language sentences.

Following Karl Lambert (Lambert 1983: 40), we distinguish between the core of a theory of predication and the specific inflection of this theory by the TTP. A we shall see, Meinongianism rejects the specific inflection but agrees with the core of the TTP. The core of the theory can be expressed via two claims:

> **(T9)** The logical form of a predicative proposition is that of the concatenation of a predicative   expression *F(x)* and an individual constant *a*, such as in: *F(a)*.
>
> **(T10)** A proposition of the form *F(a)* is true just in case *a* is *F*.

The traditional theory, however, adds a further thesis to (T9) and (T10), namely:

> **(T11)** For a sentence of the form *F(a)* to be meaningful, "*a*" must be assigned a unique existent   element of the domain.

Lambert, with the additional ascription of the origins of the standard theory to Russell and Frege, states (T11) as follows:

> There is a respect in which Meinong's theory of predication differs not one iota from either Russell's theory or Frege's (as manifested in Frege's conception of scientific language). This respect is what will be called the *core* of the traditional theory of predication. Roughly it says that the truth-value of a predication depends on whether what is said of the object specified by its singular term (or terms) is true (or false) of that

object (or objects). However, there is also an important *constraint* on the traditional theory adopted by Russell and Frege, but rejected by Meinong. It says, roughly, that the objects specified in a predication must have being. (1983: 40)

Theses (T9)-(T11) deliver a systematic notion of a predicative proposition, *i.e.*, any proposition that respects these three theses is predicative, while any proposition that fails to meet them drops out of that class. By extension, any natural language sentence that expresses a thus defined predicative proposition counts as predicative, and any natural language sentence that fails to express such a predicative proposition is not predicative.

Now that we possess a theoretical criterion for identifying predicative sentences, we can start to examine which natural language sentences are predicative. Here, however, the issue becomes quite complex. Even if we know the formal requirements established by (T9)-(T11) that a proposition must meet, the question of which natural language expressions—if any—are the equivalents of logical individual constants is not a trivial one. Furthermore, without an answer to that question, we cannot properly carve out the class of natural language sentences that express predicative propositions.

Following the cue of Meinongianism, we shall assume for the sake of the argument that at least some natural language proper names can be treated as the equivalent of logical individual constants. On the basis of that conjecture, we shall try to move forward step by step[15]. This assumption in place, we can then state that a sentence such as:

---

[15] Russell himself denied, at least for a while, that proper names could be interpreted as semantically equivalent to logical individual constants, restricting the this function to indexicals like "I" (Russell 1905). However, as our aim merely is to draw a contrast between two theories of predication, we shall ignore such complications.

(5) John is happy

is predicative, interpreting it as "Happy (John)." Let us now compare (5) with the following:

(6) All men are happy.

Is (6) also a predicative sentence? One reason that we might think so is that we can flawlessly substitute "John" for "All men" in (6) and again obtain (5). As such, one might take this syntactic observation as evidence that these two expressions play the same role, *i.e.*, they provide an object of which happiness is then predicated to the propositions expressed by (5) and (6)[16]. One might claim that this simply amounts to taking the grammatical form of (5) and (6) at face value, and this form seems to be the same in both cases: subject-copula–predicate.

Logicians, however, have long contested this assertion and denied that sentences like (6) are truly predicative. Why? In a nutshell, logicians are essentially interested in the inferential powers of propositions, and, as it turns out, (5) and (6) have very different inferential powers. We can make this more concrete as follows:

(7) Sam is a man.

Assuming that both (5) and (7) are true, we can only—by applying the valid inferential rule of conjunction introduction, *i.e.*, "$p, q \vdash p \mathbin{\&} q$", to be read as, if p is true and q is true—infer the truth of the following:

---

[16] This, of course, is simplified for the sake of the argument. There are also syntactic observation of the same kind that speaks *against* this claim. See, *inter alia*, Sainsbury 2001.

(8) John is happy, and Sam is a man.

Admittedly, inferring (8) from (5) and (7) is not tremendously interesting. However, from the fact that (6) and (7) are true, we can infer something much more substantial. Indeed, if (6) and (7) are true, we can infer the following:

(9) Sam is happy.

Hence, (5) and (6) have very different inferential powers. Moreover, these powers must be reflected in the logical form of the proposition expressed by (6).

The standard logical form of a sentence like (6) is the following quantified form:

(10) ∀(x) Man (x) —> Happy (x).

The assignment of this logical form to (6) makes it quite obvious why (9) follows from (6) and (7). Indeed, if anything that is a man is happy, and if Sam is a man, then it follows—by *modus ponens*—that Sam is happy.

The general idea at play here is that something important is revealed when we transition from, say, (6) to (10). That is, we clearly and transparently demonstrate the conditions under which (6) is true. As Mark Sainsbury puts this point:

> Formalizing a sentence or argument of a natural language in one [artificial language]… reveals something about the nature of the natural language, something that would otherwise be apt to remain hidden. (Sainsbury 2001: 339)

Coming back to (5) and (6), the process of formalization allows us to clarify their truth conditions and, consequently, their inferential powers. On the basis of these remarks, we can formulate the following thesis:

> **(T12)** For any natural language sentence S with a form of type F, there is no guarantee that the proposition expressed by S also instantiates a form of type F.

Hence, generally speaking, for each natural language sentence S, we can distinguish between what is standardly called its *grammatical* form and its *logical* form. This distinction in place, we can then fully capture the meaning of the notion of *interpreting* a natural language sentence into a formal language: An interpretation of a natural language sentence must aim to bring the logical form of the proposition it expresses to the surface, and this form can be hidden behind its grammatical form. The process of interpretation, however, brings it to light.

Therefore, what we learn from (5)-(10) is that, as a general rule, we must distinguish between the grammatical form and logical form of a natural language sentence. More locally, natural language sentences formed by means of an expression like, 'All $x…,$' do not express predicative propositions. Rather, they convey quantified propositions.

Thus far, Meinongians have little upon which to frown. Indeed, the above naive syntactic test that we proposed to identify predicative sentences, *i.e.*, replacing "all men" with "John," turned out to be too simplistic. As a matter of fact, (5)-(10) reveal that the semantic contributions of expressions like"'all men" and "John" are very different. The first are quantified expressions that, in a standard analysis of quantifiers, are second-order concepts that quantify the number of instances of some F. On the other hand, as we assumed above, the second are the equivalent of logical individual constants and refer to

particular individuals. Meinongians, however, become concerned when we consider how to treat sentences like the following:

(11) Pegasus does not exist.
(12) Pegasus has wings.
(13) Vulcan was postulated by Leverrier.

Indeed, at least *prima facie*, all of these sentences seem to share a grammatical structure that is identical to (5) and (6). Additionally, contrary to the case of (6), it seems that we should be able to readily interpret them as predicative sentences. Moreover, (11) and (13) are simply true, while (12), at least for some, is intuitively true. Using "E!(x)" to denote a first-order existence predicate permits us to interpret each of these as follows:

(14) ¬E!(Pegasus)
(15) Winged (Pegasus)
(16) Postulated-by-Leverrier (Vulcan)

The problem, however, is that according to the TTP and assuming that neither Pegasus nor Vulcan exists, (14)-(16) cannot be taken as faithfully expressing the logical form of (11)-(13). Indeed, the TTP uses (T11) as a constraint on predication. Since neither Pegasus nor Vulcan exist, despite the initial plausibility of treating (11)-(13) as predicative, following the TTP means that we must simply relinquish this intuition. Alternately, we must treat these sentences in one of the following three ways:

(a) They express a proposition, but their subject matter has nothing to do with Pegasus or Vulcan. Rather, they are concerned with arbitrarily assigned referents, such as the empty set, as proposed by Frege in the *Grundgesetze*.

(b) They fail to express any proposition, and so do not have any meaning.

(c) They express a proposition that is somewhat about Pegasus and Vulcan, although it does not have a predicative logical form.[17]

Options (a) and (b) are dead ends. Indeed, (a) is committed to the claim that all sentences formed by means of a singular term that fail to refer to something that exists are *de facto* about the same thing. However, this is implausible. First of all, (11) denies existence of Pegasus, not of the empty set—we assume here that existence is a first-order predicate, not what is expressed by "∃". Moreover, (11) and (12) are about the same thing, while (13) is about something distinct from them. Nevertheless, if they all refer to the empty set, then we simply cannot account for this fact. Hence, we should reject option (a).

Option (b) is even worse. It is a fundamental fact that such sentences do have meaning. Indeed, we must be able to deny the existence of things, and so (11) must have a meaning. Similarly, it seems impossible to deny that (12) and (13) also do have a meaning. Therefore, the standard theory cannot treat them as meaningless on pain of performing a *reductio ad absurdum* of itself. The only viable choice is thus option (c).

Option (c) amounts to treating predicative sentences formed by means of a singulars term that are about things that do not exist as expressing propositions with a logical form that substantially deviates from their grammatical form. Granted that logical sentences can take only take one of two forms, a predicative and a quantified one, meaningful sentences that cannot be treated as predicative must therefore be treated as quantified.

---

[17] For discussions of each of these views see, *inter alia*, Frege 1959 for the first; Evans 1982 for the second; and Russell 1905 for the third.

At this point, proponents of the traditional theory routinely appeal to Russell's theory of descriptions. We shall discuss that theory in more detail in Chapter 4, and so shall merely say the following for the time being: According to Russell, we can interpret all definite descriptions as quantified sentences and all proper names as truncated definite descriptions. We can endorse this result in this strong form, or we can restrict it to the case of singular terms that are about things that do not exist. As we shall only focus on the latter kind of singular term, this distinction does not matter much for the time being. We shall ignore the case of singular terms that are about things that exist and focus on those that are about things that do not. What shall matter for us is that as a result of endorsing a Russellian treatment of such terms, natural language sentences such (6)-(8) must be interpreted as having the following logical form:

(17) ¬∃(x) Pegasizes (x)

(18) ∃(x) Pegasizes (x) & Winged (x) & ∀(y) Pegasiszes (y) & Winged (y) —> x=y

(19) ∃(x) Planet (x) & Postulated-by-Leverrier (x) & ∀(y) Planet (y) & Postulated-by-Leverrier (y)   —> x=y.

In terms of the material introduced in this section, we can then drive home the following point: The TTP has important consequences for the doctrine of predicative sentences. Indeed, it carves out the class of natural language sentences that express predicative propositions in such a way that many natural language sentences that have a predicative surface grammar turn out to have a non-predicative logical form. Some might find these results to be piercing and understand them as freeing us from many philosophical difficulties. Russell put forward the treatment of negative singular existential statements like (11), for instance, as a solution to the ancient Parmenidean riddle of non-being, and some

celebrated it as such (see, *inter alia*, Quine 1953). That said, the standard theory of predication is a disaster for the IOM. Indeed, it not only affects sentences such as, "Pegasus has wings," but also sentences such as, "John is thinking about Pegasus." If the traditional theory is true, we cannot, contrary to the IOM, treat these as predicative sentences. A Meinongian theory of predication is hence committed to rejecting the standard theory. It is here that the principle of independence kicks in, and Meinongianism uses that principle to carve out of the class of natural language sentences expressing predicative propositions, thus offering an alternative to the the account of predication provided by TTP.

4.4 The Meinongian Theory of Predication

Meinong's principle of independence constitutes the core of a Meinongian theory of predication, and its endorsement comprises a key step in remolding the TTP. On the one hand, Meinongianism *endorses* both (T9) and (T10) and is thus highly structurally similar to the TTP. That said, it flatly rejects what it considers an undue restriction, namely, (T11). As a result, it delineates the boundaries of the class of natural language sentences that express predicative propositions in a very different manner than the TTP.

   The rejection of (T11) possesses two aspects. The first one concerns the possible referents of singular terms. According to Meinongianism, there are more things that those that exist, and so there is no reason to claim that the only possible referents are existent ones. The principle of indifference accounts for this possibility, and the principle of independence accounts for the claim that these objects can have properties. The second aspect concerns the rejection of classical logic as the preferred formal language for interpreting natural language sentences. Indeed, it is an axiom of classical logic that each individual

constant is assigned exactly one *existent* referent. Neil Tennant, in a typical account of how classical logic deals with singular terms that fail to be about something that exists, explained this matter:

> Our language allows the formation of terms such as 'the square root of Jupiter' or 'the empty set's wife.' Are we to regard these as denoting any objects? Our present answer is simple and evasive. We design our language so that this problem never arises. We secure every name a denotation, and we assume that every function is 'everywhere defined.' (Tennant 1990: 22)

From a Meinongian perspective, this endorsement of classical logic is unnecessary. Indeed, the principle of indifference provides us with positive reasons to reject the assumption that all individual constants must be assigned an existent referent, since all objects have their natures determined *Außersein*. As such, there is no reason to deny that an individual constant can, first of all, be assigned a non-existent referent and, secondly, that this non-existent referent can be the object of a semantic evaluation to determine whether a certain predicative expression is true of it.

According to Meinongianism, we can simply claim that a sentence such as, "Pegasus has wings," is a predicative sentence, one that results from the concatenation of a proper name that refers to a non-existent object and a predicative expression and that is true just in case the predicate "has wing" is true of that non-existent object. Of course, this does not yet tell us *how* we can evaluate such sentences and, hence, *whether* they are true or false. It does, however, allow us to treat them as predicative sentences and to then carve out the class of natural language sentences that express predicative sentences in a much less restrictive way than the TTP would imply.

The difference between the TTP and Meinong's principle of independence can be made more explicit by reflecting on the validity of some inferences. According to the standard theory, the following inference is deductively valid:

(20) *a* is F;
(21) *a* possesses a positive ontological status.

Indeed, according to the standard theory, it is a matter of logical truth that if (20) is true, then (21) is also true, because the axiomatic of classical logic forbids an individual constants from being assigned non-existent referents. Meinongianism, however, rejects the validity of this inference. Indeed, according to Meinongianism, since the *Sosein* of an object is independent of its *Sein*, no inference can be drawn from a true premise like (20) to (21).

As Kenneth Perszyk (Perszyk 1998) makes clear, however, we would do well to not read too much into this claim of the *Sosein's* independence from its *Sein*. In particular, the logical force of the independence in question does not extend beyond the fact that there is no logical implication between a predicative expression being true of an individual constant and the referent of that individual constant possessing a positive ontological status. It would be mistaken, however, to read this statement of independence as the claim that a predicative expression being true of something and its referent having a certain kind of ontological status are logically independent in the stronger sense that A is logically independent from B if, and only if, A neither entails B nor its negation. Indeed, as mentioned above, Meinong distinguishes between the nature of an object and what follows from this nature. Accordingly, it might well be the case that, when considered as a pure object, an object's instantiation of a certain property is independent from

its *Sein*. That said, as mentioned above, this does not mean that it cannot follow from the object's nature that it either cannot exist or must exist.

This point needs emphasis, as it will play an important role in the later discussion of a classical objection to Meinongianism. Indeed, the gist of this argument is that if Meinong believes that *Sosein* is independent from *Sein*, he does not believes that the opposite is true, *i.e.*, that *Sein* is independent from *Sosein*. On the contrary, his claim that a contradictory object, such as the round square, cannot exist is equivalent to the claim that we can logically infer, on the basis of an object's *Sein,* that there are some properties, *e.g.*, contradictory ones, that it cannot instantiate. In other words, if an object has *Sein*, then there is a range of properties that it cannot instantiate, namely, those that would make it impossible for it to have *Sein*. Hence, Meinong's principle of independence does not entail a theory in which anything goes. It merely demands that pure objects, *i.e.*, objects considered from the point of view of the determination of their nature *Außersein*, are such that their *Sein* is independent from their *Sosein*. However, it also entails that *Sein* is not independent from *Sosein*, as *Sein* is incompatible with certain kinds of *Sosein*.

4.5 Consequences of the Principle of Independence

Meinong's rejection of the standard theory and its replacement by the principle of independence involves a rejection of classical logic. Thus, a particular kind of non-classical logic must replace it: a free logic, which is a logic that is "free of existence assumptions with respect to its terms, singular and general" (Lambert 2002: 258).

In fact, Meinongianism requires a specific version of free logic. First of all, it must be a so-called *positive* free logic, namely, a free logic in which the concatenation of an individual constant that refers to a non-

existent object can not only be meaningful, but also true. Moreover, it needs a free logic with a neutral interpretation of quantification, *i.e.*, one that does not interpret quantification as ontologically committing. We shall call such a logic a *Meinongian free logic*.

Such a logic is traditionally modeled in terms of dual domain semantics. On the one hand, neutral quantifiers range over a non-empty outer domain that can be conceived as filled with Meinong's objects determined *Außersein* and a (possibly empty) inner domain—according to free logic, that something exists is not a matter of logical truth—that can be conceived as the domain of what exists. Absolutely unrestricted quantification is then conceived as quantification over the larger domain, and unrestricted quantifiers are *unloaded*, *i.e.*, non-ontologically committing, quantifiers. The traditional quantifiers of classical logic are then interpreted as restrictions of these unloaded quantifiers, *i.e.*, as these quantifiers only ranging over existent things. As Graham Priest explains it:

> Free logics of the kind at which we have been looking contain names for non-existent objects, but they do not allow us to quantify over them. This may be thought somewhat arbitrary, especially given the semantics. Why not allow quantifiers to range over all objects? Thus, we might add another kind of quantifier whose truth conditions are exactly the same as those in classical logic, with domain of quantification D…Let us call such quantifiers *outer quantifiers*, as opposed to the quantifiers with domain E, which are *inner quantifiers*. (Priest 2008: 295)

Such a Meinongian logic, however, needs to introduce new symbols to be able to distinguish between outer and inner quantifiers. Different authors choose different typographical conventions to do this. We shall, however, choose the following ones: We will write the unloaded particular quantifier as "$\sum$" and read it as "some." There is *no* connection between

the meaning of this sign as the summation sign in mathematics and the way we use it here. The unloaded universal quantifier, on the other hand, will be written as "√" and read as "for all." One again, there is *no* connection between the meaning of this sign as the square-root sign in mathematics and our usage.

That these quantifiers are "unloaded" means that there can be true positive quantified statements about things that do not exist, such as, "Some objects do not exist." We can interpret these as follows: (where "E!(x)" stands for a first-order existence predicate):

(22) $\sum(x)\ \neg E!(x)$

We can then treat the two quantifiers of classical logic, "∃" and "∀", as restricted quantifiers, that is, as quantifiers only ranging over the inner domain. Importantly, such a Meinongian free logic is not committed to endorsing two equally fundamental pairs of quantifiers. Indeed, just like classical logic, it is committed to only one fundamental quantifier, either "$\sum$" or "√." Assuming that one chooses "√" as the fundamental quantifier, "$\sum$" can be standardly introduced as follows:

(23) $\sqrt{(x)}\ F(x)$ *iff* $\neg\sum(x)\ \neg F(x)$

The classical logic quantifiers can then be defined as following (see Priest 2005: 14):

(24) $\exists(x)\ F(x)$ *iff* $\sum(x)\ F(x)\ \&\ E!(x)$
(25) $\forall(x)\ F(x)$ *iff* $\sqrt{(x)}\ E!(x)\ F(x)$

As we shall see below, some critics have objected to the development of such a Meinongian free logic. In particular, they have claimed that its

interpretation of quantification is contradictory. The critical part of this chapter evaluates such arguments.

4.6 Descriptive Support for the Principle of Indifference

One might well wonder, however, where Meinong thinks he can find support for his endorsement of such a principle as substantial as the principle of independence which, as we just saw, impinges not only on the theory of predication but also on the kind of logic and notion of property that one should endorse. As noted by Richard Routley, the support for this second principle is essentially descriptive. Our warrant for accepting this claim is due to the fact that there are, according to Meinongianism, many examples of the correct attribution of properties to non-existent objects. As Routley puts it:

> The Independence Thesis, that items can and do have definite properties even though nonentities, *is supported by a wide range of examples* of nonentities to which definite properties are attributed. These attributions occur when people make true statements about items, and therefore ascribe properties to them, without assuming them to exist or knowing full well that they do not exist. These examples represent counterexamples to the Ontological Assumption, unless a successful reduction of the example statements to statements about entities is produced. They therefore provide a *prima facie* case against the Ontological Assumption. (Routley 1980: 28–emphasis added)

What Routley calls the ontological assumption is classical logic's specific claim that individual constants must be assigned existent referents in order to be meaningful. According to Routley, we have descriptive grounds for rejecting this thesis. However, Routley only claims that this establishes a *prima facie* case against the ontological assumption. In other

words, according to Meinongians, we must first of all carve out the axiomatic of the theory of objects and then assess it on the basis of what follows from it. As such, the claim that there are non-existent objects could still be defeated if it ran into insoluble troubles. Moreover, as we shall see when we reach the critical part of this chapter, many authors have claimed that endorsing the principle of independence indeed creates more problems than solutions. First, however, the next section present a final Meinongian principle, the CP.

5. The Characterization Principle

The last principle that we shall discuss, the CP, both plays a role in determining the nature of the pure objects and a role in how we can know truths about them. Regarding the first point, it tells us how the nature of an object can be determined *Außersein*. Regarding the second point, it tells us how we can know which properties a non-existent possesses.

As a first pass, the principle can be formulated as follows: For any non-empty set of properties, some object possesses all of these properties. For instance, some object possesses all of the properties in the following set: {being a detective, being quick-witted, being addicted to cocaine, living in 19[th] century London}, namely, Sherlock Holmes. In other words, Meinongianism's most general domain of quantification can be roughly conceived as the function that maps the set of all the properties there are to all its possible non-empty sub-sets, as each such sub-set corresponds to an object. As Graham Priest clarifies:

> If $A(x)$ is any property, or conjunction of properties, we can characterize an object, c_A, and be guaranteed that $A(c_A)$. This is the *Characterization Principle* (CP) in its most naive form. (Priest 2005: 83)

Thus understood, the CP both plays a role in determining the nature of pure objects and has an epistemic function. First of all, it tells us how the nature of a pure object is determined *Außersein*, *i.e.*, in terms of the function just mentioned. Epistemically speaking, it tells us how we can *know* what nature a pure object possesses—namely, the one just determined by this function. This knowledge, moreover, is *a priori*. It simply follows from a knower's rightful interaction with the CP.

As we shall see, our above first pass at the CP reveals itself as too strong. It simply cannot be the case that for any non-empty set of properties *S*, there is something, existent or not, that corresponds to *S*. Indeed, if this were the case, then it would simply follow from the CP that anything exists. Indeed, for any set of properties that does not contain the existence property, we could simply add it and consequently define anything into existence. The CP, in other words, needs to be restricted, and some properties must be expelled from its range.

The problem for Meinongianism, however, is expelling the problematic properties from the CP in a way that is not purely *ad hoc*. In other words, it must restrict the CP in manner that does not progress case by case, as trouble-makers are encountered along the way. Rather, we need a systematic methodology for restricting CP. In *Über Möglichkeit und Wahrscheinlichkeit*, Meinong proposes sharply distinguishing between two kinds of properties that he calls, respectively, "*konstitutorische*" and "*außerkonstitutorische,*" standardly translated as "nuclear" and "extranuclear" properties (Meinong 1915: §25). Nuclear properties, according to Meinong, are all properties that play a role in determining an object's nature. Extranuclar properties, on the other hand, are those that do not. His claim is that the CP should be restricted to nuclear properties.

As we already know from the principle of indifference, existence does not determine an object's nature. Meinong, however, adds more properties to the list of the extranuclear properties. In *Über Möglichkeit*

und Wahrscheinlichkeit, for instance, he discusses the case of simple particulars, namely, particulars with only one property, such as colors. The shade phenomenal 'blue,' for instance, possesses just one property: being an unmixed instance of an apperance of blue. Nevertheless, Meinong remarks, this means that this shade in fact possesses two properties: being blue and being simple. The latter property, he contends, does not play a role in determining its nature. It is, contrary to the first property, extranuclear.

Another important kind of extranuclear property are intentional properties, namely, the fact that some object is being thought about. Sherlock Holmes, for instance, possesses the property of being thought about right now, as you read this sentence. That is, according to Meinong's endorsement of the IOM, thinking about Pegasus means that one stands in a relation to Pegasus, and so Pegasus possesses the relational property of being thought about. Such a property, however, does not determine an object's nature. In other words, Sherlock Holmes' nature is determined separately from whether anyone is thinking about him, and being thought about is merely one of his extra-nuclear properties.

As we shall see later on, some authors have objected to the CP, even in its restricted form. However, we shall postpone our discussion of that topic until the critical part of this chapter, which begins in the next section.

6. Objections to Meinongianism

Meinongianism has not been, to put it mildly, a popular view. Since its origin, philosophers have treated it with suspicion, and it has been the object of many objections. In this section, our goal is not to be exhaustive. Rather, we present a series of objections of different force. Some of them merely point to some of Meinongianism's commitments

without necessarily constituting an ultimate challenge to the coherence of that view. Some, however, do greater damage and seriously challenge Meinongianism's credentials as a plausible theory. Thus, on the basis of these objections, we ultimately reject that view and, consequently, the IOM.

6.1 The Nonsense Objection

The first objection we shall discuss the principle of independence in a straightforward manner. We call   it the nonsense objection. It is a common complaint, but one of its most recent iterations can be found in the work of Uriah Kriegel. Its gist is quite simple: The mere idea that non-existent objects can instantiate properties is *nonsensical*. Hence, the principle of independence cannot be correct. Here is how Kriegel puts it:

> It is sometimes claimed that there is a special class of relations, "intentional relations," that do not require the existence of the relata. Strictly speaking, however, this is nonsense. The only way it could make sense is if a two-place relation could be instantiated even where there are not two relata. Yet nobody thinks it remotely plausible that a *monadic* property could be instantiated even when there is no entity that instantiates it, *e.g.*, that squareness could be instantiated even if there are no squares. That is clearly absurd. The same absurdity attaches, I contend, to the parallel claim about relations: just as a monadic property cannot be instantiated in the absence of an instantiator, so a relation… cannot be instantiated in the absence of a relata. (Kriegel 2011: 140)

Upon reading Kriegel's quote, however, one might feel slightly uneasy. Indeed, he begins by describing the view he targets as one in which some relations do not require the existence of their relata. This perspective, he

claims, is "nonsense." He then goes on to assert that this view is identical with one in which a relation would obtain "even where there are not two relata." He then goes on to state that nobody believes that a property can be instantiated where there is no entity to do so. One might wonder, however, how the claim that some relations do not require the existence of their relata entails that a relation can obtain with no relata and, moreover, that a property can be instantiated where there is nothing to instantiate it. A Meinongian, in any case, would deny that the latter two claims follow from the first one.

According to Meinongianism, it is indeed true that some relations can obtain even if one of their relata does not exist. However, this does not mean that such a relation is *lacking* a relata. It means, rather, that one of its relata is a non-existent one and, according to Meinongianism, non-existent relata are just like any other relata. Secondly, that a property can be instantiated even if nothing instantiates it clearly is an absurd view. The problem, however, is that Meinongianism never claims such a thing. The claim that *Sosein* is independent from *Sein* does not amount to the claim that *Sosein* can be instantiated on its own, without something also being some way. It means, rather, that a non-existent particular can instantiate a property in just the same way as an existing one. In sum, the problem with Kriegel's above quote is that it clearly attacks a straw man. The view it describes indeed is absurd, but no one, or at least no Meinongian, endorses it.

We might, however, recast Kriegel's objection so that it bears on Meinongianism. Properly expressed, the criticism is not that it is absurd that a property can be instantiated by nothing but that it is nonsensical to claim that something that does not exist can instantiate properties. Is this second version of the objection any better?

As a matter of fact, it is not. First of all, some distinguished philosophers have believed, and still believe, that non-existent objects can instantiate properties (*inter alia*: Meinong 1960; Parsons 1980; Routley

1980; Priest 2005; Berto 2013). It seems a slightly uncharitable assessment of their beliefs to dub them as nonsensical. Secondly, we do continuously talk as if non-existent objects could have properties. We say things like, "Sherlock Holmes is bright" or"Pegasus has wings." Hence, the claim that such things cannot actually have properties certainly requires more of response than simply claiming that it absurd. Thirdly and finally, *instantiation*, *property*, and *particular* are all technical notions. As such, the claim that the application of these notions to the realm of non-existence is absurd certainly needs a bit more discussion than a mere statement. Indeed, we might ask, why *exactly* is it absurd? On its own, the nonsense objection does not seem to deliver any answers to this question. Let us then move on to the next objection.

6.2 The Meta-Ontological Objection

The central question of ontology amounts to the following: "What are our ontological commitments?" To answer this question, however, one must be able to answer a preliminary one, namely, "What *criterion* should we use to determine our ontological commitments?" Indeed, as Philipp Bricker puts it, "A criterion of ontological commitment is a pre-requisite for ontological inquiry" (Bricker 2014). Hence, one of the central tasks of meta-ontology is answering this preliminary question.

Contemporary meta-ontology is characterized by an "orthodox view" (Bricker 2014; see also Berto & Plebani 2015 and Takho 2015) according to which we are ontologically committed to whatever falls in the scope of the quantifiers of what we take to be our ultimately regimented, existentially quantified sentences. In other words, according to this orthodoxy, existential quantification is the meta-ontological criterion that we must use in our debates on ontological commitments. The *locus classicus* of this orthodox view is Quine's influential essay, "On What There Is," in which he writes:

> We can very easily involve ourselves in ontological commitments by saying, for example, that *there is something* (bound variable) which red houses and sunsets have in common; or that *there is something* which is a prime number larger than a million. But this is, essentially, the *only* way we can involve ourselves in ontological commitments: by our use of bound variables. (Quine 1953: 12)

This orthodox view shall interest us insofar as it constitutes the basis of an important objection against Meinongianism and, hence, against the IOM. Indeed, since the IOM is committed to the truth of thesis (i)—some intentional objects do not exist—mentioned in the introduction of this chapter, it shamelessly quantifies over non-existent things. If, however, quantification is indeed ontologically committing, then thesis (i) is simply contradictory, as it would tell us that some things both exist (implicitly, in its quantified part) and do not exist (explicitly, in its predicative part). The meta-ontological objection is thus that the IOM is, *via* its endorsement of Meinongianism, committed to a kind of existentially non-loaded quantification, *i.e.*, Meinongian quantification, and that there simply *cannot* be such a thing.

The full, explicit objection (inspired by Berto and Plebani 2015: Chap. 7) goes as follows:

1) In the logical framework of the orthodoxy, *i.e.*, first-order classical logic, that something exists is treated as the claim that something is self-identical with an element of the domain, *i.e.*, "*a* exists" is interpreted as "$\exists(x) \, x = a$;"

2) Negative existential claims, on the other hand, are interpreted as the denial of the claim that something is self-identical with an element of the domain, *i.e.*, "*a* does not exist" is interpreted as "$\neg\exists(x) \, x=a$;"

3) Accordingly, *"there are things that do not exist,"* must be interpreted as: "∃(x)¬∃(y) x=y;"

4) This, however, is the equivalent of "∃(x) x≠x," and, granted that everything is self-identical, this amounts to a contradiction;"

5) Hence, Meinongianism is false.

In this section, our goal is to evaluate the force of this objection. Before turning to that task, however, we shall begin by distinguishing between two different contexts, or ontological projects, within which such an objection can occur. After explicating this distinction, we shall find ourselves in a better position to evaluate it.

6.2.1 Descriptive *vs.* Revisionary Ontological Approaches

The most general context within which the meta-ontological objection can be formulated is that of using a formal language to clarify the content of natural language sentences. As it turns out, however, this general idea can be implemented in two very different ways, a "descriptive" and a "revisionary" one (see, *inter alia*, Kriegel 2008; Crane 2013: 29), thus giving rise to two versions of the above objection.

To begin, the descriptive approach takes the transition from a natural language sentence to its interpretation in a formal language as a move whose aim is to uncover the real structure, as opposed to the surface structure, of the propositions expressed by natural language sentences and, hence, to acquire insights into the proper meaning of natural language sentences. To use an example mentioned above, if one utters, "John did it for Sam's sake," and takes this sentence to be true, we might wonder what exactly is the content of this assertion and to what exactly one is committed by endorsing its truth. The transition from the natural language sentence to a formal interpretation is supposed to help

us to answer that question by getting a handle on the underlying logical structure of the proposition expressed by that sentence.

We shall characterize the descriptive approach in terms of the interaction between three things: (a) natural language sentences, (b) their interpretations in a formal language, and (c) their ontological commitments. The first are the initial data, the second are the output of the interpretation process, and the discovery of the third is what the transition from (a) to (b) should ease. By the interaction between these three things, we mean that each of (a)-(c) is liable to place *some* theoretical pressure on the whole process. That is, according to the descriptive approach, sorting out our ontological commitments is a process that results from the interaction of three different forces: native speakers' linguistic intuitions at level (a), the logical form as determined at level (b), and, finally, ontological commitments determined at the level (c). In other words, it is not ruled out *a priori* that some reasoning operating at level (a) can pressure level (b) or (c), and, to some extent, the other way around.

We can nevertheless say that these interactions are governed by a specific norm that we shall call the *norm of adequacy*. Though the label is ours, the idea of such a norm is inspired by Crane, who writes:

> The descriptive approach is concerned to get as much of our natural language right as possible, and it is evaluated against the linguistic judgments ('intuitions') of native speakers. (Crane 2013: 30)

This quote has two parts. In the first part, Crane clearly outlines the aim of the descriptive approach, namely, providing an adequate picture of the content of natural language sentences. In the second part, he hints at a standard or norm for evaluating the descriptive approach: We must evaluate these descriptions in the light of the linguistic judgments, or intuitions, of the language's native speakers.

We contend that the second part of the quote is implying that the transition from a natural language sentence to its formal interpretation should be evaluated on the basis of the judgments of native speakers of that language. In that sense, we might still disregard an interpretation of a sentence that presents an advantage regarding the evaluation of its ontological commitments if it leads us too far from the linguistic intuitions of native speakers. In other words, linguistic intuitions can put legitimate pressure on the cogency of one's chosen interpretation and, hence, on one's resulting ontological commitments. In that sense, the above-mentioned elements (a)-(c) indeed interact, but this interaction is governed by a overarching norm that reflects the aim of the descriptive approach: providing descriptions of natural language sentences and, then, a clearer picture of *their* ontological commitments.

The revisionary approach, on the other hand, conceives of the relation between a natural language sentence and its formal interpretation in a very different manner. Of course, it still distinguishes between the above (a)-(c). However, from the revisionist point of view, the aim is not to use the second to describe the content of the first in order to get clear on the third in such a way that the overall process is accountable to the intuitions of native speakers. Instead, the objective is to use the second in order to revise the first in a manner that guarantees its agreement with ontological commitments that have been determined on independent grounds.

Within such an approach, the norms at play in evaluating the transition from natural to formal language are very different from the those at play in the descriptive approach. In such a context, there is nothing like the above-mentioned norm of adequacy. On the contrary, the dialectic runs in the other direction; it is independently determined ontological commitments that place pressure on the rest. If a certain formal interpretation is better aligned with some independently determined ontological commitment, we should select it even if it greatly

deviates from the meaning of the original natural language sentence, as assessed on the basis of native speakers' linguistic intuitions—hence, the label of revision. As Berto and Plebani put it:

> …if the aim of the ontologist is not to clarify the real ontological import of speakers' ordinary talk, but to reform it so that it better sticks to what is out there, she may not need her paraphrases to do what we ordinarily expect from good translations. (Berto & Plebani 2015: 40)

What are the sources of the norms at play in the revisionary approach? This is a question with many ramification, but we might simply point to two main sources of that view and of the norms that govern them. The first source is the formal language, or family of formal languages, standardly endorsed by revisionism as the canonical idiom in which to conduct ontological enquiries. The second source is a certain general view as to what shape the best theory of reality should take and which data should allow us to determine our ontological commitments, independently of a concern for natural language and native speakers' intuitions.

As an example of the first kind of reason, we can point to Quine's tireless defense of first-order classical logic as the canonical idiom in which to pursue ontological debates. Non-exhaustively, his main grounds are that this formal language is purely extensional, *i.e.*, all its connectors are truth functional, its co-referential terms can be substituted *salva veritate*, all of its referential positions can be subjected to the application of existential generalization, and, last but not least, the limitation of that language to first-order quantification guarantees that, coupled with a nominalist account of the values of predicates, one is not committed to quantification over anything else than concrete particulars.

As an example of the second kind of reason, we can again look to Quine, who is well-known for his tireless defense of full-blown

naturalism, defined as the view that the language of fundamental physics has complete primacy over any other modes of expressions and that a central philosophical task is either to provide translations of philosophical theories into that language or explain them away. From such a perspective, our ontological commitments should not be determined by philosophy but by science itself, and by physics, in particular. In that spirit, the task of ontology is not to determine our ontological commitments from its own authority but to provide a methodology that can help us to clarify as much as possible the ontological commitments of a given theory. Whether we are actually committed to whatever the translation in the canonical idiom reveals we would be committed to if the theory were true ultimately depends on the scientific community's decision regarding the status of the theory in question.

This distinction between these two ways of conceiving of the ontological battlefield, a descriptive approach and a revisionary one, should now allow us to differentiate between two versions of the meta-ontological objection: a descriptive one and a revisionary one. Let us begin with the revisionary version. Revisionism is, first and foremost, a methodological doctrine. Indeed, it claims that we have independent, methodological reasons for endorsing the claim that quantification is ontologically committing—that is, that a statement that existentially quantifies over some x entails an ontological commitment towards that x. Indeed, as Quine indicated, the goal of revisionism is to provide a neutral ground—a so-called canonical idiom—from which we can assesses the ontological commitment of any kind of theory.

From that point of view, the notion that quantification is not ontologically committing does not make much sense. The premise of revisionism is that we should use a language with loaded quantifiers, because such a language provides us with a sound meta-ontological methodology. As such, if revisionists felt pressured by Meinongianism,

they could always answer that they use a stipulated form of quantification for specific means—and, moreover, that we should all follow them down that path, because such a condition bears important philosophical fruits. As such, the problem with the revisionist objection is not that it is uninteresting. Rather, the issue is that it is formulated from within a debate that occupies a rather insular position with respect to the discussion of Meinongianism. As such, we shall simply leave the revisionist objection aside and exclusively focus on the descriptive version.

6.2.2 The Descriptive Approach and the Meta-Ontological Objection

In this section, we focus on how we should understand the meta-ontological objection from a descriptivist perspective. From the revisionist point of view, the objection against Meinongianism amounted to the claim that Meinongian quantification contravenes the *stipulated* meaning of the existential quantifier. The descriptive approach, however, does not involve itself in the business of stipulation. Instead, as Crane describes, it is, "concerned to get as much of our natural language right as possible" (Crane 2013: 30). From that perspective, the meta-ontological objection cannot be that Meinongian quantification contravenes the stipulated meaning of the existential quantifier. Rather, the critique is that the existential quantifier turns out to be ontologically committing as a matter of its *natural* meaning in the vernacular. The gist of the objection is therefore that Meinongianism cannot be true, because sentences that quantify over things that do not exist are contradictory as a matter of their natural meaning.

Let us make this claim more precise by means of the following example:

(26) Some things do not exist.

As such, (26) is a sentence of the vernacular and, moreover, one that can be both meaningful and true. Indeed, imagine a tense discussion between a young child and her parents over the non-existence of Santa Claus. Furious, the child refuses to believe that Santa Claus does not exist. The parents, on the verge of a nervous breakdown, utter, "Get over it; some things don't exist." On the basis of this observation, one might well wonder what exactly is the orthodoxy's charge against Meinongianism. Indeed, one might conclude from such an example that whatever theory of quantification one endorses, one ought to be able to account for the truth of (26).

As a matter of fact, however, Meinongianism and the orthodoxy mean very different things by their putative recognition of the truth of (26). Indeed, according to Meinongianism, (26) is not only true but also *literally* true. That is, according to Meinongianism, (26) must be taken as meaning that among those things that "some" ranges over, we can discriminate between those that possess the property of existence and those that do not. The orthodoxy, on the other hand, claims that (26) can only be *non-literally* true. In other words, it is true not by virtue of the fact that its Meinongian interpretation is true but because a paraphrase of (26) that does not quantify over anything non-existent is true. According to the orthodoxy, (26) merely constitutes a sloppy way of saying something like, "There is no such thing that fits the description commonly attributed to 'Santa Claus'."

The orthodoxy's case against the literal truth of (26) rests on the idea that, read as such, (26) is contradictory. In other words, it is equivalent to the claim that some things that exist do not exist, since "some things" already *means* "some existent things." This, however, constitutes a flat contradiction, on par with a sentence such as:

(27) John is human and not-human.

The whole objection can be summarized under the form of a *modus tollens* as follows:

1) Meinongianism is committed to the claim that at least some sentences that quantify over non-existent things are literally true.
2) If at least some sentences that quantify over non-existent things are literally true, then the existential quantifier is not ontologically committing.
3) The existential quantifier is, as a matter of its meaning, ontologically committing.
4) Hence, Meinongianism is false.

The key premise of the argument is, of course, premise 3). We shall now discuss it in detail.

6.2.3 Quantification and Existence

As the name suggests, quantification is the ability to speak about *quantities* rather than particulars elements of a domain of discourse. When used unrestrictedly, universal quantifiers like the natural language "all" or "everything," and their formal equivalent "∀" permit us to speak about everything that is in the domain of quantification. When overtly or covertly restricted, on the other hand, they allow us to speak about a sub-set of the domain. Existential quantifiers, on the other hand, such as the natural language "some" or "a handful of" and their formal equivalent "∃," permit us to speak about some elements of a certain sub-set of the domain. If true, "Some x are F," for instance, tells us that there is a non-empty sub-set of the domain, namely, F.

This first function of quantifiers we might call their *basic function*, as all conceptions of quantification should embrace it. Allowing us to speak about quantities is simply the purpose for which quantifiers are designed. According to the orthodoxy, however, quantifiers do something more. That is, according to the orthodoxy, quantifiers not only allow us to speak about quantities, but also ontologically commit us to whatever makes quantified statements true. As Tim Crane puts it:

> If the standard approach is right, then we cannot think of 'some *Fs*' as picking out a collection of things independently of whether they exist. And so we cannot go on to predicate existence of some of them but not of others. This is because 'some' already introduces, implies, or otherwise contains the idea of existence. A defender of the standard view might say that this is the reason that the symbol used to represent 'some' in the predicate calculus ('∃') is called the *existential* quantifier. (Crane 2013: 31)

One might wonder, however, why exactly the orthodoxy thinks that the two functions we just mentioned always come together. In the vernacular, it is pretty clear that not all uses of a quantifier are existentially committing. The grammatically well-formed sentence, "Some characters in the Bible exist and some do not," is a perfectly suitable example of a natural language sentence in which quantification and existence diverge. According to the orthodoxy, however, such sentences only *seem* to quantify over non-existents. Why should we think that these two functions, speaking about quantities and expressing ontological commitments, always go hand in hand? In answer, we shall point to *three* main reasons that may be advanced in favor of that claim.

The first reason is that an account of quantification that allows for literal quantification over non-existents must countenance the claim that non-existent things can fall within the scope of quantifiers. Hence, such an account of quantification must presuppose that there are things that

do not exist. This first claim, however, can hardly serve as an objection against a Meinongian theory of quantification. Indeed, saying that such a theory of quantification commits to non-existent objects is a mere statement of facts. What would be needed is an argument against the existence of such non-existent objects. This claim, however, does not provide any such argument.

On that score, as Priest (Priest 2008) and Berto (Berto 2013) point out, it is quite ironic that in what is universally considered as the *locus classicus* of the view that quantification is ontologically committing, namely, Quine's "On What There Is" (Quine 1953), one cannot find a single positive argument in favor of the view. Rather, Quine proceeds by elimination, deflating proper names and variables of any form of ontological commitment before asserting that quantification commits us ontologically. He merely *assumes* that the existential quantifier is ontologically committing and does not provide any additional support in favor of his conjecture. At best, Quine makes clear that he does not believe that there are things that do not exist and, hence, that quantification is always quantification of what exists. This claim, however, is not about quantification itself but, rather, about *domains* of quantification. As Berto writes:

> [In 'On What There Is'] there is no argument positively supporting the thesis that existential commitment is expressed by quantification: Quine *assumes* that a domain of quantification can encompass only existing things. (Berto 2013: 31)

In other words, if Berto is right, then of the two Quinian claims—(a) quantification is ontologically committing, and (b) domains contain only things that exist—it is the latter that wears the trousers. That is, Quinian quantification is ontologically committing solely because domains only contain existing things. However, from a Meinongian point of view, it is

precisely this conception of domains that is challenged. Thus, we can simply leave this first reason behind us.

Peter Van Inwagen has advanced a second, more substantial, kind of reason. The gist of Van Inwagen's claim is that the mere notion of Meinongian quantification is problematic, because it amounts to, in Berto's words, "something like an analytic falsehood" (Berto 2013: 73). That is, the assertion that, "There are things that do not have any kind of being," amounts to a falsehood by virtue of the meaning of the terms employed. One cannot meaningfully say that "there *are*" things before denying that these things "have any kind of being." Here is how Van Inwagen puts it:

> Meinongianism entails that there are things that participate in *neither* mode of being, things that have no being of any sort; but if there *are* such things, they obviously have being. For a thing to have *being* is for there to *be* such a thing as it; what else could being be? (Van Inwagen 2008: 39)

Let us, first of all, clarify a background assumption in Van Inwagen's objection, namely, the claim that one cannot appeal to a distinction between kinds of being to deflate the charge that Meinongianism is contradictory. Indeed, reflecting on sentences such as, "There are things that do not exist," one might simply claim that this contains no contradiction on the grounds that we have to contrast what it is for something to *have being* from what it is for it to *exist*. A thing that exists certainly has being, but the converse need not be true: One thing may have being but fail to exist.

Such a rejoinder assumes a substantial thesis—that there are kinds of being and that existence only constitutes *one* kind of being among others. One main intuitive ground for the doctrine of the plurivocity of

being is that different things exist in different ways: Substance are wholly present at each time of their existence, events unfold over time, properties are abstract, and so on. This, however constitutes an unwarranted reason for supporting the claim that we can distinguish between different kinds of being. Indeed, as Van Inwagen puts it, "many philosophers distinguish between a thing's being and its nature" (Van Inwagen 2009: 476), and this distinction should suffice to account for any metaphysical difference between, say, substances and events. Everything that exists does so in the very same sense, even though things that exist may possess different natures. To be wholly present at each time and to unfold over time, for instance, does not amount to have different kinds of being but to having different kinds of natures.

We shall simply grant this background assumption to Van Inwagen. As a matter of fact, the debate over the plurivocity of being and the debate over Meinongian quantification turn out to be *orthogonal*. Indeed, some Meinongians accept the thesis of the univocity of being and nonetheless defend a Meinongian conception of quantification (see, *inter alia*, Berto and Plebani 2015: 107). Thus, from a Meinongian perspective, if there is a problem with Van Inwagen's objection, it is not that he denies that we can distinguish between kinds of being.

This first rejoinder set aside, let us now come back to Van Inwagen's argument against Meinongian quantification proper. We shall propose reconstructing it as follows:

1) Meinongianism is committed to the literal truth of, "There is a thing that does not possess any kind of being."
2) The meaning of the expression "a thing has being" is given by, or is equivalent to, the meaning of the expression "there is a thing that…"

3) If 2) is true, then any true sentence formed by means of an expression of the form "there is…" entails that some thing has being.

4) Hence, Meinongianism is committed to a contradiction.

The key premise of this argument is premise 2), as 1) merely provides information about Meinongianism, and 3) is supposed to follow from 2). Is it, however, plausible? Van Inwagen makes it sound as if it is a rather straightforward point. However, 2) contains a very substantial notion—and one that, moreover, is poorly supported by evidence. Let us see how a proponent of Meinongianism might answer Van Inwagen's objection.

We shall discuss two complementary lines of answer to Van Inwagen's objection discussed in the literature on quantification and non-existence. The first one is that the use of the verb "to be" is *accidental* to quantification (Berto 2013; Crane 2013). That is, there are many different forms of quantification in the vernacular, and not all of them use the verb "to be." The second one is that the verb "to be" is ambiguous and so not all expressions of the form "there is…" or "there are…" mean the same thing as "to exist" (Berto 2013).

Natural language quantifiers are much more rich than the generic "∃" of first-order classical logic. Indeed, "there is…" and "there are…" are not the only means of speaking about quantities. In addition, expressions like "some" or "a handful of" are fully-fledged quantified expressions, even though they do not make use of the verb "to be." Also, one would need a supplementary argument to support the claim that all of these quantified expressions are nonetheless synonymous with "there is…" and "there are…," despite the fact that they do not overtly contain the verb "to be." Van Inwagen's argument, however, is insufficient for establishing that claim.

Note, moreover, that even if such an argument existed, it would not be sufficient to substantiate Van Inwagen's claim. Indeed, one would

also need to establish two further claims: first of all, that "to be" is synonymous with "to exist" and, secondly, that all such quantified expressions are generically translatable by "∃." As we shall now see, the existence of an argument seems very unlikely.

Indeed, the verb "to be" is ambiguous between "absolute" and "non-absolute" uses (Berto 2013: 75). In an absolute use, "to be" is equivalent with "to exist." This usage is very well attested in technical contexts, like philosophy, but not that common in the vernacular. Indeed, philosophers often say things such as, "There are numbers" or "There is a God," or even, "God is," but these constructions are not very common in the daily routines of competent English speakers. On the contrary, in the vernacular, non-absolute constructions are the most common, that is, constructions in which the verb "to be" is—just like the verb "to have"—conjoined with adjectives to form predicates. One says, for instance, "John is late" or "John is Swiss," just has one says, "John has measles" or "John has big dreams."

Now, clearly enough, in an absolute use, "to be" *is* synonymous with existence. What, however, about the non-absolute uses? Is "to be" also synonymous with "to exist" in such cases? As Berto points out, this is rather unlikely. If we take, for instance, a sentence such as, "There are two trucks in the garage," which is felicitous, and replace it with "There exist two trucks in the garage," we seemingly end up with something much less felicitous (Berto 2013: 75). As such, then, there does not seem to be a straightforward route from the claim that *absolute* uses of "to be" are synonymous with "to exist" to the stronger one that *all* uses of "to be" are synonymous with "to exist."

In answer, one could ask whether it is at least the case that even if absolute and non-absolute uses cannot be synonymous, we can say that non-absolute uses *entail* existence, *i.e.*, that when one asserts that "John is late," one does not *mean* that John exists but says something that at least

entails that John exists. As a matter of fact, this question is precisely the one at stake within the debate over Meinongianism. Hence, one cannot simply claim that there is such an entailment in the absence of an argument, as this would simply beg the question against Meinongianism. Thus, the orthodoxy must offer us a substantial argument, but, as we shall now see, the prospects for such an argument do not look very promising.

Indeed, linguists distinguish between two uses of expressions like "there is…" and "there are…." On the one hand, they can have an *existential* function, in which they assert the existence of something and, on the other hand, they can have a *locative* function, in which they merely introduce elements of discourse. Now, for Van Inwagen's argument to work, the relations between these two uses would need to be such that: either existential and locative uses of such constructions would be co-extensional or, at least, locative constructions would constitute a sub-set of existential ones, *i.e.*, all locative constructions would be existential ones, even though some existential constructions might not be locative ones.

As a matter of linguistic analysis, however, the contrary is the case: Existential constructions constitute a sub-set of locative constructions. In other words, existential constructions are a sub-type of the introduction of elements of discourse—the introduction of *existing* elements of discourse. The descriptive version of the meta-ontological objection is hence doomed to remain unsupported by the linguistic intuitions of native speakers.

Frederike Moltmann argues for this point in her paper, "The Semantics of Existence" (Moltmann 2013). She notices that the locative uses of quantificational expressions are much broader than their existential uses. That is, in the vernacular, there are uses of "there is…" and "there are…" that explicitly do not convey any ontological commitment. As Moltmann remarks:

…there-sentences can quantify over past, merely possible, and merely intentional objects, *objects that the predicate exist (or other existence predicates) could not be true of*. This is particularly clear with [noun phrases] containing intensional or intentional adjectival or relative clause modifiers, as in the following, possibly true sentences:

(13)   a. There are possible buildings that do not actually exist.

                b. There are imaginary buildings that do not exist.

                c. There are historical buildings that no longer exist.

(14)   a. There are buildings built in the past that no longer exist.

                b. There are buildings I might have built that do not exist.

                c. There are buildings John thought of that do not exist.

(M o l t m a n n 2013:7–emphasis added)

This allows us to deliver the final blow to Van Inwagen's argument: He cannot simply ground his claim in the *meaning* of the expressions "there is…" and "there are…." That there is an *is* in "there is" is hence not sufficient for arguing that any such expression is ontologically committing. That said, as Berto points out, all of this does not mean that there are no uses of "there is…" and "there are…" that may *partially* support Van Inwagen's claim (Berto 2013: 77). That is, there may well be cases where "there is…" and "there are…" entail existence. In light of the above consideration, however, such uses must be considered as restricted, rather than absolute, uses of quantifiers.

These restrictions, moreover, need not be explicit. That is, in order to quantify over things that exist, one may not need to say something like, "Of the things that exist, some of them are merely contingent." On the contrary, it is often sufficient to assert, "There are things that are

merely contingent," where it is routinely taken up by the context that the quantifier "There are" is restricted to things that exist. As such, the Meinongian diagnosis of Van Inwagen's mistake is that he treats what are, in fact, restricted uses as absolute ones and hence misinterprets the relation between existential and locative constructions. As Berto concludes, "That there is an 'is' in 'there is' should not lead us to conclude that, in all cases, we ascribe… being in the absolute sense or existence, to the things we quantify over" (Berto 2013: 75). To conclude, Van Inwagen's argument has no grip on Meinongian quantification.

Finally, a third reason that could be put forward as an objection against the idea of Meinongian quantification is that we need the notion that quantification is ontologically committing, because we possess no other means of expressing ontological commitments. In particular, there is no hope of carving out a *bona fide* first-order existence predicate that does not appeal in part to quantification, like the standard first-order predicate of first-order classical logic, "E!(x) = df. $\exists y(y=x)$." Here, however, we shall simply rely on examples from the literature of such *bona fide* existence predicates (see, *e.g.*, McGinn 2001; Berto 2013; Crane 2013).

To summarize our discussion, the meta-ontological objection, when understood as a descriptive one, does not seem to have strong support. As such, we can conclude that if Meinongianism poses a problem for the IOM, it is not by virtue of its theory of quantification. In the remainder of this chapter, we shall point to three unwelcome consequences of Meinongianism: (i) its endorsement of a restriction of the law of non-contradiction, (ii) the primitive character of the distinction between nuclear and non-nuclear properties, and (iii) the strongly unintuitive character of its doctrine of nature.

6.4 Meinongianism and the Law of Non-Contradiction

The next objection we shall consider is that Meinongianism entails a *rejection* of the law of non-contradiction, *i.e.*, that it cannot be the case that both *p* and ¬*p* are true or, in the material mode, that a particular instantiates both a property and its negation. This law is standardly qualified as one of the basic laws of thought, and Aristotle himself, in the *Metaphysics* Γ, famously called it "the most certain of all principles." The charge is that Meinongianism is bound to reject it.

The objection centers on the CP. As we saw above, this principle needs to be restricted in order to avoid absurd results, such as the consequence that everything can be defined into existence. Meinong's answer was to distinguish between two kinds of properties, nuclear and non-nuclear ones, and to claim that the CP is restricted to nuclear properties. The gist of the present objection is that this restriction is by no means sufficient for Meinongianism to avoid other absurdities, such as the rejection of the law of non-contradiction.

Indeed, for any nuclear property, such as being round, one can imagine a set of properties *S* such that *S* contains this property and its negation, *e.g., S*={being round, ¬being round}. Moreover, it follows from the CP that there is an object that instantiates these two contradictory properties. That there is such an object, however, contradicts the law of non-contradiction. Assuming that the law of non-contradiction indeed constitutes a basic law of thought, Meinongianism contravenes it, greatly threatening its plausibility.

Meinong's own answer to this objection is pretty simple: He bites the bullet. That is, he recognizes that there are objects whose nature contravenes the law of non-contradiction. This, one might argue, just makes it even worse for Meinongianism. Indeed, not only is the view contradictory but, on top of that, its proponents explicitly recognize that

charge. They cannot even plead ignorance, as they *explicitly* reject the law of non-contradiction.

At this point, however, one needs to be cautious when assessing Meinong's reaction to the charge of infringing on the law of non-contradiction. Indeed, he is well-aware that the rejection of the principle of non-contradiction is problematic but claims that his own infringement on the law of non-contradiction is innocuous. He does not reject it but simply proposes *restricting* it. Indeed, he claims that his rejection of the law of non-contradiction holds *only* for impossible objects. As such, then, Meinongianism's infringement on the law of non-contradiction should not come as a surprise. It merely follows from the claim that some non-existent objects are such that they possess a nature that precludes their existence.

Recall our above discussion of the principle of independence. We claimed that the meaning of Meinong's slogan "*Sosein* is independent from *Sein*" means that it does not follow from the truth of sentence of the form "*a* is *F*" that *a* exists. We also made clear that Meinong did not believe the converse of this principle, *i.e.*, that *Sein* is independent from *Sosein*. Indeed, as he explicated, some objects possess a nature that prevent them from existing. It is to these objects, Meinong claims, that the law of non-contradiction does not apply. As such, when we are talking and reasoning about ordinary objects, the law of non-contradiction holds for Meinong, just as it holds for anyone else.

Logically speaking, the worse consequence of this rejection of the law of non-contradiction is that some principles of inference that were considered valid turn out to be invalid. Graham Priest mentions, for instance, that the disjunctive syllogism would turn out to be invalid if the law of non-contradiction were rejected (Priest 2006: 110). According to disjunctive syllogism, p v q, $\neg p \vdash q$. But if we reject the law of non-contradiction, then p and $\neg p$ can both be true and so one could not rely on it to eliminate the disjunction in the first premise. Importantly,

however, Meinong's infringement of the principle of non-contradiction does not entail that he has to refrain from using such logical principles. Indeed, since he merely wishes to restrict, and not reject, the law of non-contradiction, a principle like disjunctive syllogism can still be validly applied when reasoning over everything but impossible objects. As such, then, Meinong's restriction of the law of non-contradiction merely amounts to the claim that it does not constitute a valid principle to think about impossible objects, something that, given the nature of these objects, one might have guessed right from the start. About this issue, John Findlay writes the following:

> We speak of the *world* of *Außersein*, but in reality the objects which have no being do not constitute a world. They are a chaos of incoherent fragments… From another point of view *Außersein* is incapable of scientific treatment because of its excessive richness. In the case of the actual world we can always ask whether a certain object is comprised in it or not; the question is interesting, because some things are excluded from it. The realm of *Außersein*, however, has no such exclusiveness; every possibility or impossibility is comprised in it, and this fact silences a multitude of questions. (Findlay 1963: 56-57, quoted in Lambert 1983: 58-59)

How should we evaluate this objection? If formulated as the claim that Meinong is bound to a rejection of the law of non-contradiction *simpliciter*, then it seems too strong. Meinong is not bound to a *rejection* of the law of non-contradiction but, rather, to a *restriction* of it, something that one may have anticipated, granted that impossible objects do, by definition, break that law. Hence, the objection must be reformulated as the one that Meinong is bound to a restriction of the law of non-contradiction. This restriction, however, would be truly problematic if one of its consequences were something like the complete rejection of

the disjunctive syllogism as a valid principle of reasoning over possible objects. Meinong, however, is not bound to such a rejection.

As such, then, one might simply reformulate this objection as follows: *If* one wants to endorse Meinongianism, one must be ready to live with such a restriction of the law of non-contradiction, a feature that some might find distressing. Meinongians, on the other hand, might be ready to live with it.

6.5 Nuclear and Non-Nuclear Properties

A restriction of the law of non-contradiction is not the only problem looming around the CP. Indeed, in this section, we discuss a further potentially problematic consequence of the CP, namely, the primitive character of the distinction between nuclear and non-nuclear properties.

The distinction, recall, amounts to the following. A naive formulation of the CP runs into the problem that everything can be defined into existence. Meinong's answer to this problem is to define two kinds of properties, those to which the CP applies and those to which it does not. The first ones are called nuclear properties (*konstitutorische*) and the second ones extranuclear properties (*Außerkonstitutorische*). Existence, Meinong claims, is an extranuclear property, which purportedly neutralizes the problem that everything can be defined into existence.

What, however, is an extranuclear property? Meinong's characterization of this notion is rather thin: Every property such that it creates difficulties for the CP must be deemed extranuclear. However, is there something *more* informative one can say about this distinction? As Tobias Rosefeldt remarks, this does not seem to the case:

A further disadvantage of Meinongianism is, according to me, the fact that the distinction between nuclear and extra-nuclear properties cannot be further explained philosophically. The distinction is introduced, because the conjunction of the assumption…that existence is a discriminating property…and the naive characterization principle leads to absurd consequences. The problem is that we do not possess another understanding of the notion of extra-nuclear property than the one of a property that creates difficulties for the conjunction of the characterization principle and [the assumption that existence is a discriminating property]. (Rosefeldt unpublished: 37 – own translation[18])

The objection here is not that the distinction between nuclear and extranuclear properties leads to problematic consequences. The problem, rather, is that the distinction admittedly plays an important role for Meinongianism and, as such, should be liable to thorough philosophical explanation. Meinong, however, does not provide much clarification beyond describing the *role* that the distinction should play, *i.e.*, preventing the CP from becoming an absurdity. That said, one might regard such an elaboration of the distinction as slightly too thin to play a foundational role in the elaboration of the axiomatic of a Meinongian theory of objects. Again, the objection might not be lethal for Meinongianism, as

---

18 "Ein weiterer Nachteil der Meinongianischen Konzeption ist in meinen Augen die Tat- sache, daß die Unterscheidung zwischen konstitutorischen und außerkonstitutorischen Eigenschaften keiner tiefergehenden philosophischen Erklärung zugänglich ist. Die Unterscheidung wird eingeführt, weil die Annahme (M6) – Existenz ist eine diskriminierende Eigenschaft – und das naive Komprehensionsprinzip zusammengenom- men zu absurden Konsequenzen führen würden. Das Problem dabei ist, daß man von einer außerkonstitutorischen Eigenschaft eigentlich keinen anderen Begriff hat als den Begriff von einer Eigenschaft, für die das naive Komprehensionsprinzip zusam- men mit (M6) zu Schwierigkeiten führen würde." (Rosefeldt 2006: 37)

Meinongians might be perfectly content to bite this second bullet. It is, however, another question as to whether we should really follow them in that direction.

6.6 Meinongianism and the Objection from Natures

In the last two sections, we presented two important consequences of Meinongianism: its restriction of the law of non-contradiction and the primitive character of the distinction between nuclear and extranuclear properties. In this section, we discuss a third such consequence, this time pertaining to the notion of nature, which Tim Crane proposes in his recent *The Objects of Thought* (Crane 2013).

According to Meinongianism, there are many things that do not exist, or that even cannot exist, but that are, apart from their lack of a positive ontological status, pretty much like anything that exists. That is, the golden mountain might not exist, but it is still a mountain made of gold, and Pegasus might not exist, but he is still a horse. That is, in conformity with the principle of indifference, anything can have its nature determined *Außersein*, and in conformity with the CP, this nature can be constituted of any set of nuclear properties. Hence, Pegasus and any existing horse can share the very same nature, namely, being a horse.

This might sound alright until the moment one pauses to reflect on what exactly is constitutive of a horse's nature. One might conceive that, through an unlikely cosmic event, a bunch of molecules materialize as a horse. As such, then, one might be willing to accept that being born from other horses is not a constitutive condition on being a horse. That said, what about other properties, such as occupying space and time, having inner organs, normally having four legs, and so on? Would anything lacking any of these properties still count as horse?

It is highly intuitive to believe that anything that is lacking these properties *could not* be a horse on the same count as any existing horse. However, Pegasus or any non-existent horse lacks any of these properties. Hence, the thought goes, it cannot be true that non-existent objects really share a nature with existent ones.

Meinongianism, however, is likely to answer that this objection is misguided. Indeed, Pegasus and other existent horses *do* possess properties such as occupying space-time, having inner organs, and having four legs. It is just that they occupy a non-existent space-time, have non-existent inner organs, and normally have four non-existent legs.

The problem with this Meinongian rejoinder is that it somewhat misses the original intuition behind the objection. The original idea is that non-existent horses cannot be horses on the same count as existent ones insofar as being a horse is an *existence-entailing* property, that is, it entails having a series of properties that nothing that fails to exist could instantiate. Meinong, we saw earlier, distinguishes between a thing's nature and what follows from this nature. Any object, including a necessary one, does not have existence as a constituent of its nature. A necessary object, according to Meinong, is not an object that possesses the property of existing as constitutive of its nature. It is, rather, an object whose existence necessarily follows from its nature as determined *Außersein*. The problem pointed to here, however, is that we greatly struggle in making sense of this distinction between a thing's nature and what follows from this nature. Indeed, quite often, a thing's nature is directly constituted by properties that are existence-entailing in the sense that one cannot instantiates this property and not exist. As Tim Crane puts it:

> The general point here should be obvious: there are general empirical or metaphysical conditions that objects normally have to meet if certain predications are going to be true of them. What these conditions are is a

matter of the nature of the individual things in question. But in many cases, it is required by something's having the property in question that if they have that property, there are other properties they must have, and many of these properties are existence-entailing. This is not because of anything about the logic of the sentences we use to ascribe these properties; but because of the natures of the things to which the properties are correctly ascribed. (Crane 2013: 62-63)

At this point, the distinction between Kant's conception of the principle of independence and Meinong's understanding of that principle becomes particularly crucial. According to Kant, the concept of something is independent from its existence. However, the concept of a horse is not itself a horse. Rather, it is something that can be satisfied by a horse. Meinong's trouble is that he believes that something can be a horse and yet not fulfill the most basic conditions of being a horse.

In his famous critique of Meinong, Russell points to what he calls "that feeling of reality that ought to be preserved even in the most abstract studies" (Russell 1919: 169). He claims that by believing that some things do not exist, Meinong loses touch with this sense of reality. Crane, however, proposes reformulating Russell's critique in terms of the idea that some natures are existence-entailing. The problem with Meinongianism, hence, is that it entails that *all* natures are not existence-entailing. As Crane puts it:

I would like to interpret Russell's phrase in terms of the understanding of having properties just offered. Objects have natures. What their natures are is a matter of empirical or metaphysical study. But having some of these natures requires that those objects exist. It is in the nature of horses, planets, golden things, living things (and so on) to exist. Non-existent things do not have what it takes to have the properties of these things. Non-existent objects cannot have properties like being a horse,

being golden, being a detective, and nor can they stand in relations like killing…The feeling for reality, then, is a feeling about what is required by the natures of things. Someone who lacked the feeling for reality might mistakenly think that something could be a planet even if it never actually orbited anything. But this is a mistake: it belongs to the natures of certain things that they exist. Meinong's theory of objects, I think, is based on this mistake. (Crane 2013: 63-64)

Two remarks are in order about this last objection. First of all, once again, a Meinongian might just reject it on the grounds that it uses a mistaken conception of natures. As such, then, the objection might not turn out to be lethal. It is, however, quite strong. Indeed, Crane's account of the conditions that pertain to certain natures sounds plausible. Moreover, a Meinongian willing to defend herself against Crane's criticism would have to defend some particularly strange theses, such as that there is a non-existent space-time. In fact, one might start to wonder what such theses even *mean*.

Secondly, Crane's account is a charge against the principle of indifference and the CP. That is, it attacks the idea that all natures that conform to the CP are determined *Außersein*. Importantly, however, Crane's critique is compatible with a restricted endorsement of the principle of independence. That is, *if* some properties turn out to be non-existence entailing, there is no reason to deny that, for these properties at least, putative non-existent objects could have them. This opens up the possibility of a very deflated version of Meinongianism, which we shall discuss in Chapter 6.

Why not turn to such a discussion right now? In addition to the objections already discussed, a further extremely influential objection against Meinongianism is not merely that it has troubling consequences but also that such a strange philosophical theory is simply unnecessary. Whatever task the theory of objects was supposed to help accomplish—

in particular, providing an adequate theory of intentionality—can be achieved independently of any Meinongian inkling. We shall now turn to this objection by discussing two alternative models of intentionality in Chapters 4 and 5.

7. Conclusion

As we saw in the previous chapter, the IOM is committed to the truth of the following two theses:

(i)   Some intentional objects do not exist.

(ii) Non-existent intentional objects can be the relata of the relation of "being about."

In this chapter, we discussed a particular theory that aims to account for their truth, namely, Meinongianism. As we saw, it intends to account for these two theses in a substantial way, *i.e.*, in terms of their being mind-independent non-existent objects that have their nature determined *Außersein* and to which intentional experiences can stand in the relation of "being about."

We then discussed five main objections against Meinongianism: the nonsense objection, the meta-ontological objection, the objection from the law of non-contradiction, the primitive character of the distinction between nuclear and non-nuclear properties, and, last but not the least, the objection from natures. As we saw, none of them constitutes a lethal threat for Meinongianism. Taken together, however, they start to pile up and to constitute a rather heavy bullet for Meinongians to bite. The cost of being able to account for (i) and (ii) might start to look prohibitive. If we add that many have found Meinongianism to be, in fact, unnecessary, then one might consider that

it is simply time to move on to non-Meinongian alternative theories of intentionality.

That said, this does not necessarily mean that all of the fundamental ideas of Meinongianism are hopeless. Rather, it means that if they are to have any plausibility, they must take a very restricted form, one that is unlikely, however, to be capable of accounting for (i) and (ii) as they stand. We shall return to this issue in Chapter 6. In the meantime, the next chapter turns to the IOM's most important relational contender, the content model.

Chapter 4: The Content Model

> It might seem odd to suggest that experiences might not have content, since it is difficult to find any major figure in contemporary discussion who does not characterize the intentionality of experiences in terms of their possession of content. Pick up any work on perception, intentionality, or the philosophy of mind, that has been written in the last several decades, and the chances are that it is bristling with 'contents'.
>
> (Hopp 2011: 8)

1. Introduction

We started our inquiry with the ideas of both the fundamental intentional schema (FIS) and a literal interpretation of that schema, as offered by the intentional object model (IOM). We then stumbled onto the main problem faced by the IOM, namely non-existence, and claimed that the IOM should be conjoined with Meinongianism. As we saw, however, Meinongianism faces some concerns that undermine its plausibility. In this chapter, we hence move on to an alternative relational model of intentionality, the content model (CM). We discuss a *simple* version of the CM, as we contend that the problems it faces are rather basic.

One may try to formulate the key idea behind the CM in this manner. The IOM considers the following two ideas to be two faces of the same coin: (i) the underlying metaphysical nature of intentionality is

relational and (ii) intentionality is aboutness. Indeed, the IOM asserts that intentionality is a relation and that this relation is nothing but the "being about" relation. In contrast, the core characteristic of the CM is that it proposes treating these two ideas separately.

On the one hand, the underlying metaphysical nature of intentionality is indeed relational; however, first, this relation is not the one of "being about" and, second, its relata are not intentional objects but rather intentional contents. Nonetheless, intentional contents are not what intentional experiences are about. As such, the underlying metaphysical nature of intentionality is then indeed relational, although this relation is not one of "being about"; it is instead one of standing in a *non*-intentional relation to a content.

On the other hand, according to the CM, intentionality's central feature is indeed aboutness—but not aboutness as a relation. The alternative idea is that the entities one stands in relation to, namely intentional contents, have certain non-relational properties from which we can cash out the general intuition that lies behind aboutness. According to the CM, being the bearer of an intentional state is thus a kind of two-steps process. First, one must stand in a relation to a content; second, this content must have some properties that account for the fact that this state is about something. The strategy of the CM is hence to divide and conquer whereas the IOM treats matters as a single conceptual package.

This chapter proceeds as follows. We start by introducing the notions of *content* and *representation* before applying them to intentional experiences. We then discuss how the CM is supposed to account for the six features of intentionality introduced in Chapter 1, focusing especially on aboutness and non-existence. Finally, we close by proposing two fundamental objections to the CM.

2. The Idea of Content

2.1 Intentional *vs* Representational Content

Upon realizing that the hypothesis that intentional experiences can be relations to non-existent intentional objects is problematic, a somewhat natural fallback is to keep the IOM's relation of "being about" steady while trying to change its target whenever non-existence kicks in. This can be done in three main ways: first, by claiming that intentional experiences that are about things that do not exist, *e.g.*, Pegasus, are not about particulars but rather about *properties*, such as the property of being Pegasus; second, by claiming that such intentional experiences are not about non-existent concrete particulars but rather about existent *abstract* objects; and third, by claiming that such intentional experiences are not about non-existent concrete particulars but instead about *representations* of such particulars.

All of these ideas are non-starters, and for the very same reason. Indeed, we know that if John thinks about Pegasus or desires to ride Pegasus, he neither thinks about nor wishes to ride a property, an abstract object, or a representation. Indeed, whatever these things could mean, they are just plainly absurd as an analysis of such experiences' intentional character. There is simply no plausible way to go down any of these roads. One cannot get around the problem that theories of intentionality face when non-existence kicks in by simply redirecting the mind to something other than, *e.g.*, a non-existent concrete particular. Focusing on the second and third of the above possibilities, Hopp nicely summarizes this point as follows:

> If Timmy hopes that Rudolph the Red-Nosed Reindeer lands on his roof tonight, what existing object could plausibly be taken to be the object of his hope? One is tempted to start with the obvious culprit, Rudolph, and claim that he is really an idea in the child's mind or an

abstract object of some sort. Both claims are quite unbelievable. First, nothing "in" Timmy's mind, nor anything abstract, could possibly qualify as what he is thinking about when he thinks about Rudolph; as confused as he might be, he does not hope that one of his ideas or an abstract object will land on his roof tonight. What he hopes will land on his roof is a reindeer that is identical with Rudolph, and this thing does not exist, either in his head, in the actual world, or in an abstract realm. (Hopp 2011: 14)

These ideas set aside, one might nonetheless wonder whether there is not something intrinsically correct about bringing the notion of representation into the game. Indeed, Pegasus himself does not exist, but there are many representations of Pegasus that do exist and these representations manage, in one some way or another, to direct us on to Pegasus. This fact, after all, is rather startling—and, one might add, inspiring. Indeed, since representations of Pegasus manage to direct us on to Pegasus, is there not another way to use the notion of representation—*i.e.*, *not* as what the mind is directed upon—that would somewhat allow us to explain the intentional character of an experience? However, what other roles could representations play? What about the idea that intentional experiences *themselves* are representations of some kind?

This idea has been enormously influential in recent philosophy of mind. In an encyclopedia entry on intentionality, Alex Byrne symptomatically writes the following:

The moral of the paradoxes of intentionality is that thinking of intentionality in terms of "the intentional relation" is a bad idea. A better way involves drawing a distinction between the *representational content* of a mental state (or some other thing that has intentionality) and the objects (if any) the mental state is about. (Byrne 2006: 407)

However, what do we mean by *representation* and *representational content*, and how can intentional experiences themselves be conceived as representations? These are the key questions we target. That said, we actually begin with the obvious. When reading the above quotation from Byrne, one could well claim to not see much of a difference between the IOM and the CM. Indeed, after all, the IOM also distinguishes between the content and object of an intentional experience. How can Byrne's sketch then stand as an alternative to the IOM? All of the differences are obviously revealed in Byrne's italicized text. Indeed, there is an important difference between (i) distinguishing between the content and object of an intentional state in the manner of the IOM and (ii) claiming that the contents of intentional experiences are *representational* contents that themselves can account for the intentional nature of intentional experiences. However, where does the difference lie?

First, the IOM distinguishes between intentional objects and intentional contents but does not view this distinction as having anything to do with the task of explaining the intentional character of intentional experiences, *i.e.*, their having aboutness. Indeed, as we saw, the IOM claims that for an intentional experience to have aboutness just means that it has an intentional object. The notion of content is introduced at a later stage to account not for aboutness in itself but, rather, for the feature we called aspectuality. However, the matter is viewed differently within the CM, which asserts that the intentional nature of an experience must be accounted for in terms of its having a content alone, independently of its having an intentional object.

Second, as Byrne clarifies, the CM uses a notion of content that is *representational*. We detail exactly what is meant by this qualification below, but a key issue is that if a content *C* is representational, it can, unlike intentional content of the IOM, represent something in a way that is completely independent of there being something that it represents. For

instance, a map of Tolkien's Shire can represent the Shire in a way that is completely independent from anything like the Shire itself either existing or non-existing in a Meinongian way. Crucially, this is not the case with the notion of content introduced by the IOM, which is precisely defined in terms of the notion of intentional object, *i.e.*, a content as conceived by the IOM amounts to a way an intentional object is given. According to the IOM, the notion of intentional object hence cannot be dropped without thereby also losing grip on the notions of intentional content and, also, intentionality itself.

2.2 Introducing Representation

Now that these preliminary remarks have been provided, we can turn to the next question: What is a representation? Let us tackle this by discussing the example of a map of London that one is handed as leaving a hotel. We claim that the map constitutes an indisputable instance of a representation of London. In the most general terms, a map can be characterized as a pattern of ink marks that stand in some relation with each other in such a way that they can be said to *represent* the geography of some real place, that is, to represent the respective positions and distances that separate a series of real geographical locations. For instance, one point on the map represents Big Ben whereas another represents Westminster Abbey. Knowing that the top of the map represents north (and where it is), the map-reader is in a position to know the respective positions of these two landmarks and the distance that separates them. Knowing where one is on the map also enables one to reach them.

That said, in virtue of what is a map a representation? Most basically, for something to be a representation it is essential that it stands in a *semantic* relation with what it represents. Indeed, as David Pitt puts it: "a representation is an object with semantic properties" (Pitt 2012). It is

this possession of semantic properties that turns a piece of paper with scattered ink marks into a map. That is, a map contains information about the way things are in the real world and, as such, can be evaluated semantically as being either *accurate* or *inaccurate*. It is accurate if the points on the map can be said to accurately model the relations between a series of real geographical locations and inaccurate if not. We refer to the feature of the map that can be semantically evaluated as its *representational content*.

Moving beyond maps, all representations share a general commonality: they possess a representational content that is semantically evaluable. As shown below, different kinds of representations can have different kinds of representational contents and different kinds of representational contents can possess different kinds of semantic properties. Nonetheless, deep down all representations possess a representational content, to be characterized as follows:

> **Representational content:** For any representation *R*, its representational content constitutes an element or feature of that representation that can be semantically evaluated, *i.e.*, can be deemed correct or incorrect, accurate or inaccurate, true or false, and so forth.

Representations, however, are complex objects and their having a representational content only constitutes one of their features. Distinguishing among some of these features will help us make stand out what representational content is. With respect to any representation *R*, the following features can be distinguished:

a) Its *possessing* a semantically evaluable content;
b) The way it *encodes* this content;

c) The way this encoding is *realized* in a certain medium of representation; and

d) The manner it *possesses* this content.

A map of London again serves as a useful example. Such a map possesses some semantically evaluable content about London, but this content is encoded in a particular way: As a map is essentially a drawing, its content can be said to essentially encoded in a *pictorial* way. It is difficult to determine exactly how a content can be encoded pictorially, but we do not pursue this matter here. It is likely that different kinds of pictorial content encode information in different manners, and our task is not to provide a detailed account of each option. However, in the case of a map it is plausible that the notion of mapping plays an important role in the way the representational content of a map is encoded. That is, an accurate map is a map that stands in a one-to-one relation of mapping with the portion of the world (including not only particulars but also some relations between them) it aims to represent. As such, at least one way for a pictorial representation to encode its content is to be an image that maps one-to-one what it aims to represent.

Imagine now a very long text that presents a detailed description of the geography of London as it is represented on a map of that city. Such a text might plausibly be said to represent the geography of London. However, sentences are not drawings; that is, we do not claim that the sentence "Big Ben is located near the Thames" represents the geographical location of Big Ben by virtue of physical properties that somewhat map one-to-one the real relation between Big Ben and the Thames, as would be true for the kind of pictorial representational contents that characterize maps. On the contrary, such a sentence can be said to represent the real relation between Big Ben and the Thames only *symbolically*, that is, by means of symbols that are endowed with meanings which when combined in a sentence indicate that Big Ben is located near

the Thames. As such, we must then distinguish between two things: the representational content of a representation and the *particular way* in which it encodes this representation, *e.g.*, in a pictorial, symbolic, or other manner.

Next, distinguishing between different ways in which a content can be encoded, *e.g.*, pictorially or symbolically, might be finessed further by distinguishing not only the kind of a representation but also the particular *medium* in which a representation of a certain kind is realized. For instance, compare a painting and a photograph of London. While both might be said to be pictorial representations of London, paintings and photographs are very different kinds of things. As such we might distinguish between the way in which a representational content is encoded, *e.g.*, pictorially, and the *medium* into which it is encoded, *i.e.*, its *representational vehicle*. A painting and a photograph of London can hence be said to be two pictorial representations of London with different kinds of representational vehicles.

Finally, we indicated above that a map of London has been assigned a certain function, namely representing a portion of the world. This results from the imposition of a certain content to the map. However, not all representations are such that their having a content derives from their having been assigned a content by something distinct from them. Indeed, if this were the case, we would be confronted by a regress that threatens the possibility of their being representations in the first place. As such, we must thus distinguish between representations that represent *originally* or *intrinsically* and those that represent only *derivatively* (Searle 1992; Haugeland 1998a). Original representation occurs only if the fact that a representation is representing something is not explained in terms of the representational properties of something else.

2.3 Experiences as Representations

With these basic elements in place, let us return to the topic of intentional experiences. We have introduced the notion of representation as well as some key ideas that revolve around it. The idea embodied in Byrne's above quotation is then that intentional experiences form just another kind of representation, alongside *inter alia* maps, words, sentences, paintings, and pictures. Accordingly, intentional experiences are complex things in at least just the same way any kinds of representations are.

     At the most basic level, the idea is that experiences are a kind of representational vehicle, alongside maps, paintings, and pictures. However, experiences stand out among the possible different kinds of representational vehicles. Indeed, their having representational content in the first place is quite unique in the fact that they have representational content intrinsically or originally. That is, unlike artefacts endowed with representational content by means of this content having been imposed on them, experiences do have such contents in a way that is not explained in terms of the representational properties of something else. But in what way is the content of an intentional experience encoded? We voluntarily do not address this question here. Indeed, it is very likely that different kinds of mental representations encode their contents in different ways. Abstract thought might be purely symbolic while some kinds of mental imagery might be more pictorial. Accounting for this diversity would bring us too far.

     However, intentional experiences are not exhausted by having a certain representational content. Indeed, for one to undergo an intentional experience is for one to undergo an experience of a particular kind, *i.e.*, for one to judge, desire, hope, or fear something. The "something" part is supposed to be covered by the notion of representational content, but the CM still needs to account for the fact

that one is always intentionally directed upon something in a certain manner. Accordingly, for the CM, for an experience to be intentional is for an intentional subject to take a certain *intentional attitude* with respect to what a certain representational content represents. For instance, to desire that *p* is to desire that what is represented by *p* obtains. Similarly, to judge that *p* is to judge that what is represented by *p* is the case. And so forth for each possible kind of intentional attitude. In sum, the CM thus contends that for something to be about something essentially is to entertain a certain intentional attitude with respect to something that is represented by a certain representational content. As Jaegwon Kim summarizes the issue:

> You hope that it will be warmer tomorrow, and I believe that it will. But Mary doubts it and hopes that she is right. Here we have various "intentional" (or "content-bearing" or "content-carrying") states: your *hoping* that it will be warmer tomorrow, my *believing*, and Mary's *doubting*, that it will be so. All of these states, though they are states of different persons and involve different *attitudes* (believing, hoping, and doubting), have the same *content*: the proposition that it will be warmer tomorrow, expressed by the embedded sentence "it will be warmer tomorrow." This content *represents* a certain state of affairs, its being warmer tomorrow. (Kim 2006: 269)

Now that these fundamentals elements of the CM have been introduced, let us compare how the CM and IOM respectively describe intentional experiences. According to the IOM, intentional experiences are essentially characterized by the presence of intentional objects to the mind. That is, for something to be an intentional experience is for that experience to be an instance of a determinable relation of "being about" directed towards an intentional object. Accordingly, to be the bearer of an intentional experience is for one to direct one's mind upon an

intentional object in terms of a determinable relation, *i.e.*, one can *judge* that a certain intentional object obtains, or *desire* a certain intentional object, or merely *think* about it. For the IOM, all such are instances of the mind standing in the "being about" relation to some intentional object.

According to the CM, this picture of intentionality is deeply misleading. Indeed, the CM asserts that for an experience to be intentional is not for it to stand in a determinable relation of "being about" to an intentional object; rather, it is a matter of entertaining intentional attitudes toward ways the world is or can be, as represented by so-called representational contents. Let us take the following intentional experience as an example:

(1) John hopes that Federer retires.

According to the IOM, for John to hope that Federer retires is a matter for John, or John's experience, to be standing in the relation of "being about" with an intentional object—presumably, the non-obtaining fact that Federer retires—in a certain way, that is, in a way that John hopes that this fact obtains.

The IOM therefore analyzes the intentional nature of (1) as follows. "John" designates the intentional subject, that is, the person undergoing the intentional experience; "to hope" indicates the particular way in which John is mentally directed on an intentional object; and finally the that-clause "that Federer retires" plays a double role: it states both the intentional object of John's hope, *i.e.*, what it is about, and the content of John's hope, *i.e.*, the way John in which is intentionally directed upon his intentional object.

The CM provides a very different description of this example. In particular, the CM does not claim that for John to hope that Federer retires is for John to stand in some relation toward an intentional object.

On the contrary, the CM contends that for John to hope that Federer retires is for John's experience to have a certain representational content, namely *that Federer retires*. Moreover, this representational content is not what John's intentional experience is about: John's experience is about some fact, not some representational content. According to the CM, John's intentional experience is about what this content *represents*, that is, that Federer retires.

The CM therefore analyzes the intentional nature of (1) as follows. "John" designates an intentional subject, that is, the person undergoing an intentional experience; "to hope" represents the particular attitude that the experience is an instance of, and finally, the that-clause "that Federer retires" denotes the experience's representational content. To hope that Federer retires is thus to hope that what the representational content represents obtains at some point.

In the next section, we see how the CM is supposed to be able to account for the six main features of intentionality introduced in Chapter 1.

3. The CM and the Six Main Features of Aboutness
3.1 Aboutness

The first task of a theory of intentionality should be to provide us with an account of the property of intentional experiences we have called "aboutness." In this section, we focus in more details on the way in which the CM tries to account for this property. We argue that the CM defends two claims with respect to aboutness: first, that the notion of aboutness should be accounted for in terms of the notion of representational content; and second, that aboutness thus understood is conceptually derivative of another feature of intentionality, namely semantic normativity. We discuss the first point in this sub-section and

the second one in the subsequent discussion of semantic normativity itself.

The CM's key tenet is that the feature we called "aboutness" can be elucidated in terms of the idea that intentional experiences have representational content. As should be obvious, however, the idea cannot be that the very notion of aboutness employed by the IOM can be defined alternatively in terms of the notion of representational content. Indeed, the notion of aboutness used within the IOM was tailor-made insofar as, within that framework, for an experience to be about something merely means for it to stand in a relation of "being about" with some intentional object. As such whatever notion of aboutness the CM delivers, it must be one that is different from the notion advocated by the IOM.

One way to understand the CM's strategy with respect to the way it chooses to cash out aboutness is as follows: a general intuition lies behind the fundamental intentional schema and this general intuition can be philosophically accounted for in more than one way. The IOM cashes it out in terms of the notion of intentional object, but this interpretation is at best contingent and at worse mistaken as it saddles the theory of intentionality with the intractable problem of accounting for relations to non-existent objects. There is hence room for an alternative interpretation of this general intuition.

What is this general intuition? It is, roughly, that intentional experiences are essentially concerned with the world in a way that nothing non-intentional is. This familiar idea can be recalled by means of a comparison. John weighs 80 kg. This non-intentional state is a *constituent of* the world but is in no sense *concerned with* the world. As far as we can see, it is not concerned with anything. It just is what it is, namely a state of weighing 80 kg. Now take John's judgment that Obama is American. This intentional experience is *also* a constituent of the world,

but it displays something more. Indeed, unlike John's state of weighing 80 kg, it is *in addition* concerned with the world.

The IOM cashes this general intuition out in terms of the IOM schema. For an intentional experience to be concerned with the world is for that experience to take the world, or rather a portion of it, as an intentional object. However, according to the CM the problem of non-existence as faced by the IOM must be taken as a good sign that this conception of aboutness is mistaken and that the general intuition must be cashed out in an alternative manner, *i.e.*, by means of the notion of representational content.

The CM hence asserts that for an intentional experience to be concerned with the world is for it to have a representational content endowed with semantic properties. That is, like any other kind of representation, each intentional experience possesses a certain representational content C such that there are conditions under which C can be said to be semantically evaluable, *i.e.*, true, veridical, or whatever other semantic property is required. The conditions under which C can be true or alternatively what C can be true of constitute what C is about. In other words, C determines a state of the world S such that C is true just in case S obtains. By being the bearer of an intentional experience that tokens a certain intentional content, an intentional subject is hence concerned with the world in the sense of tokening a certain content that is true just in case some conditions obtain. According to the CM, this is a sufficient gloss on the notion of general intuition that lies behind aboutness.

As demonstrated later when we discuss the CM's take on non-existence, it is important to distinguish between the claim that a representational content *determines* the conditions under which it can be positively semantically evaluated and the *status* of these conditions. Indeed, at least for some way to conceive of these conditions, they can be determined independently of being fulfilled. As such, an intentional

experience with a certain content *C* can still be said to be about what would fulfill the conditions determined by *C* even if these conditions are not fulfilled.

We can accordingly saddle the CM with an interpretation of the FIS that involves two steps. First, the word "about" is eliminated in favor of the notion of representation; second, the notion of representation is cashed out in terms of the notions of representational content. The following schema is the result:

> **The CM schema**: "*x* is about *y*" amounts to "*x* represents *y*"; and *x* can be said to represent *y* just in case there is a representational content *C* that either is, or could be, true of *y*.

As Bradon-Mitchell and Jackson summarize the issue:

> If you desire that it rain soon, your desire is about what the proposition expressed by the sentence 'It will rain soon' is about: namely, rain in the near future. Moreover, its 'aboutness' does not depend on the existence of what it is about. There may be no rain in the near future; nevertheless, that is what you desire. Your desire counts as being about future rain in that only future rain would make the proposition that you desire true, and so satisfy your desire. Again, some believe in the Devil but that is not a matter of standing in the belief relation to the Devil; there is no such thing to be related to. Rather, they believe that the Devil exists, and their belief counts as being 'about' the Devil only in the sense of that the proposition that they believe could not be true unless the Devil exists. (Bradon-Mitchell and Jackson 1996: 162)

From the IOM's perspective, what is certainly the most remarkable feature of the CM schema is the complete disappearance of the notion of intentional object. Indeed, according to the CM, the notion of

intentional object plays absolutely no role in an experience being intentional. The notion of content is instead the sole determining factor. This is the now familiar sense in which most contemporary theorists of intentionality claim that something is intentional only if it is "content-bearing" or "content-carrying", as illustrated by Kim in the passage quoted above.

That being said, there is no implication that the CM has no use at all for a notion such as the one of intentional object. On the one hand, the CM claims to be able to discharge the intuitive sense of aboutness encapsulated in the FIS in terms of the above CM schema. However, intentional experiences conceived in that manner can also be said to be about something in another, more demanding, sense. Take John's judgment that Obama is American. According to the CM, this judgment is about Obama in the sense that it has the following propositional content: *Obama is American.* As this proposition contains a succesfully referential compotent, "Obama", this judgment can be said to be a member of an important sub-set of intentional experiences that not only has aboutness in the above sense of having a content but also in the sense of being *referentially successful.* That is, this proposition and the intentional experience that has it as content both manage to refer to Obama and, since this proposition is true, the fact that he is American. A second sense of aboutness can then be defined as the feature that propositions and subjects entertaining them possess when they are referentially successful.

The implications for the CM's analysis of aboutness are as follows. The members of a sub-class of intentional experiences, namely those whose content contains a successful referential propositional component, have aboutness in two senses: first in the sense of having a content and, second, in the sense of being referentially successful. On the other hand, the members of another sub-class of intentional

experiences, namely experiences that do not contain a successful referential component, have aboutness only in the first sense.

John Searle aptly captures the essential difference between these two notions of aboutness in the following passage:

> Part of the difficulty here derives from "about," which has both an extensional and an intensional-with-an-s-reading. In one sense (the intensional-with-an-s), the statement or belief that the King of France is bald is about the kind of France, but in that sense it does not follow that there is some object which they are about. In another sense (the extensional) there is no object which they are about because there is no King of France. (Searle 1983: 17)

The treatment of aboutness in terms of both the notion of representational content and the distinction between the two notions of aboutness has important consequences for the application of existential generalization (EG) to intentional experiences as analyzed by the CM. Indeed, the CM asserts that EG applies only to the members of the first sub-set of experiences, namely those that have a successful referential component as part of their content. Take the following example:

(2) John judges that Obama is American.

This sentence is a successful instance of a judgment. That is, its propositional content is true and, crucially, "Obama" is referentially successful. From (2) we can legitimately infer the following:

(3) There is someone about whom John judges that he is American, namely Obama.

Compare this with the following:

(4) John judges that Pegasus is a horse.

In the case of (4), since Pegasus does not exist and hence "Pegasus" is not referential, we are *not* allowed to infer the following:

(5) There is something about which John judges that it is a horse, namely Pegasus.

In plain words, this means that the CM deems that contrary to the case of (5), in (4) there is nothing John is thinking about in the sense of aboutness captured by the application of EG that was at the center of the IOM. If proof is still needed, this shows that the notion of aboutness that pertains to the FIS captured by the IOM and the CM are indeed different.

However, the CM contends that what matters is that both (2) and (4) meet the conditions for another application of EG that is taken to be sufficient to capture the intuitive sense of aboutness that lies behind the fundamental intentional schema. That is, according to the CM, we can uniformly infer the following from (2) and (4):

(6) There is something that John is judging, namely that Obama is American.

(7) There is something that John is judging, namely that Pegasus is a horse.

In light of this, the CM thus describes the situation of John judging that Pegasus is a horse as follows. By virtue of having a representational content to the effect that Pegasus is a horse, in (6) John's experience is engaging with the world in a way that is sufficient to describe it as being about something. However, in another sense of "about," since in (2)

John's experience is not referentially successful, then it fails to be about anything.

As the below discussion of non-existence reveals, proponents of the CM take this distinction between the two senses in which we can read "about" as an advantage of their account. Indeed, by isolating a notion of aboutness from any way the world must be like, it manages to evince the IOM's fate to postulate a realm of non-existent intentional objects that intentional experiences can be about. However, as shown below, it remains unclear whether this feature really speaks in favor of the CM.

3.2 Aspectuality

It should be recalled that aspectuality is the feature of intentionality that explains why, *e.g.*, we cannot infer (9) from (8) despite the fact that Phosphorus = Hesperus:

> (8) John believes that Phosphorus is a star.
> (9) John believes that Hesperus is a star.

How can the CM account for aspectuality? From the CM's perspective, the issue concerns the conditions of individuation of representational contents. That is, from this viewpoint the likeliest explanation of the invalidity of the inference from (8) to (9) is that the representational contents of (8) and (9) turn out to be different, such that John can entertain an attitude with respect to one content without thereby also entertaining the same attitude toward the other.

However, the problem is that these two contents, namely *Phosphorus is a star* and *Hesperus is a star*, possess the same truth-conditions, *i.e.*, they are both true if, and only if, what both "Phosphorus" and "Hespherus" refer to—Venus—is a star.

Consequently, if we individuate representational contents in terms of their truth-conditions alone, *Phosphorus is a star* and *Hesperus is a star* turn out to constitute the same representational content despite superficial differences. The conclusion to draw from this is that the easiest route for the CM to account for aspectuality is to endorse the claim that representational contents are not individuated in terms of their truth-conditions.

The question of how exactly the CM should then individuate representational contents is a very complex one that we cannot tackle here. We later expand somewhat on the determination of truth-conditions for representational contents when we discuss non-existence, but a full review and discussion of the many options would require a dissertation of its own.[19]

3.3 Semantic Normativity

Semantic normativity is the feature in virtue of which at least some intentional experiences can be said to be semantically evaluable. For instance, John's judgment that Obama is American can be said to be true or false. However, the question is in virtue of what can intentional experiences be said to be thus semantically evaluable? According to the CM, for an experience to be intentional is for it to have a representational content; moreover, for something to be a representational content is for it to possess semantic properties. From the perspective of the CM, a constitutive account of intentional experiences is hence one that must be couched right from the start in such semantic terms.

---

[19] The literature about this issue has now taken gigantic proportions. A helpful collection is Matthew Davidson's *On Sense and Direct Reference* (Davidson 2007).

In other words, the CM asserts that intentional experiences can instantiate the feature we call semantic normativity because they are representations, as defined above in semantic terms. As such it is then in the very nature of intentional experiences to have such semantic properties. At least at the conceptual level, we thus reach a rock-bottom level: intentional experiences can be semantically evaluable because, according to the CM, we understand what it means for an experience to be intentional in terms of its possession of a representational content that is endowed with semantic properties.

That being said, there is an important question for the CM, namely whether the definition of representational contents in terms of semantic properties constitutes a rock-bottom description *tout court* or, alternatively, whether we can provide an account of representational contents in non-semantic, *e.g.*, causal or causal-teleological, terms. For instance, Fodor writes the following:

> I suppose that sooner or later the physicists will complete the catalogue they've been compiling of the ultimate and irreducible properties of things. When they do, the likes of spin, charm, and charge will perhaps appear upon their list. But aboutness surely won't; intentionality simply doesn't go that deep. It's hard to see, in face of this consideration, how one can be a Realist about intentionality without also being, to some extent or other, a Reductionist. If the semantic and the intentional are real properties of things, it must be in virtue of their identity with (or maybe of their supervenience on?) properties that are themselves neither intentional nor semantic. If aboutness is real, it must be really something else. (Fodor 1987: 97)

However, it is important to note that this debate is to be located at the engineering, and not constitutive, level of a theory of intentionality. As

such, with respect to the policy introduced in Chapter 1, we simply ignore it.

3.4 Intentional Identity

Intentional identity is that feature of intentionality by virtue of which numerically distinct intentional subjects can be said to undergo intentional experiences that are the same thing. As we saw, however, the CM distinguishes between two notions of aboutness; as such, intentional identity can be taken to mean two different things. On the one hand, it can simply mean that the two intentional subjects undergo numerically distinct experiences with the same content. According to the CM, this feature can be analyzed by saying that the two subjects stand in the "having content" relation to the same proposition. That is, if both John and Sam judge that Obama is American, they can be said to be intentionally directed toward the same thing in the sense of having numerically distinct experiences that are relations to the same proposition.

That being said, in a case such as this one, John and Sam can also be said to be intentionally directed toward the same thing in the sense of both judging something about the fact that Obama is American (which reflects the second sense of aboutness distinguished above). However, this second sense of intentional identity only characterizes this sub-class of intentional experiences that can be said to possess a referentially successful representational content.

In a case in which, say, John and Sam both think about Pegasus, since they are not really thinking about anything, the question of whether their two intentional experiences are about the same thing in the second sense of aboutness distinguished above does not arise. However, we can still characterize John and Sam as sharing a same kind of intentional experience, namely a Pegasus experience, which can be

identified with having a certain content, that is a Pegasus content. Such kinds of content are addressed further in our later discussion of non-existence.

3.5 Generality

The CM can deal with generality very elegantly. It should be recalled that generality is the feature by which an intentional experience can be said to be about something, although nothing in particular. For the CM, this simply means that an intentional experience has a content with semantic properties that correspond to those of a non-uniquely restricted existentially quantified sentence. For instance, if John desires to have a beer, the CM views this as equivalent to John entertaining the attitude "desire" toward a certain proposition, namely that he has a beer. Moreover, this proposition would be true just in case John is handed a beer—any beer. This experience can thus be said to about something, although nothing in particular.

3.6 Non-Existence

As faced by the IOM, the problem of non-existence originates from the IOM's interpretation of the FIS in terms of a relation to an intentional object. Coupled with the claim that some intentional experiences can be about things that do not exist, this yields the result that intentional experiences that are about things that do not exist must be analyzed as relations to non-existent intentional objects. However, this not only commits the IOM to a realm of non-existing entities, but also to the claim that relations can obtain between denizens of distinct ontological realms: existence and non-existence. We saw in Chapter 3 that there are good reasons to stay away from these claims.

The CM does not offer a solution to this particular problem; that is, it does not offer an answer to the question: "How can there be relations to things that do not exist?" The development of the CM instead allows its proponents to find a route to non-existence that does not automatically get us trapped in the problem of non-existence as faced by the IOM. As we have seen, the CM proposes to cash out the intuitive notion of aboutness captured by the FIS in terms that are independent of the notion of intentional object. That is, the CM first constructs intentionality as a non-intentional relation to intentional contents and, second, cashes out the intuitive notion of aboutness non-relationally in terms of the semantic properties of contents.

We have seen that the CM also defines a second, more demanding, notion of aboutness that is defined similarly to the notion of intentional object used by the CM and that can be captured by successfully applying EG. That is, an intentional experience can be about something in the stronger sense of having an intentional object if its propositional content is successfully referential. However, once again, the first notion of aboutness, which is defined entirely in terms of the notion of content and of its properties, is taken to be independent of referential success. That is, a content that is not referential still possesses semantic properties and can determine what is is about, even though it is not about anything in the second (stronger) sense. To sum up, according to the CM, an intentional experience can be about something that does not exist in the sense that is has a representational content but no intentional object.

As such, the CM can then be taken to offer a *dissolution* of the problem of existence as faced by the IOM. For the CM, there is no question as to how intentional experiences can stand in relation to things that do not exist. The CM views this as a pseudo-question that does not need to be asked. Its fundamental notion of aboutness, the one taken to cash out the intuitive sense of aboutness captured by the FIS, has been

divorced from the notion of intentional object. Whereas the IOM looks for a philosophical theory, the CM offers a therapy.

One might then wish to ask how well the CM fairs as a therapy for the problem of non-existence. With a slight nod in the direction of Carnap's famous distinction (Carnap 1956), we propose distinguishing two ways of understanding this question, namely an *internal* and an *external.*

The internal question can be put as follows:

Internal question: Is the therapy offered by the CM internally consistent?

In contrast, the external question can be put as follows:

External question: Is a therapy of the kind offered by the CM the right solution to the problem posed by the feature of non-existence?

Two remarks are in order about this distinction. First, as laid out by Carnap, the distinction between internal and external questions is supposed to emphasize internal questions to the detriment of external ones. However, this does not correspond with our version of the distinction.

According to Carnap, we must distinguish internal questions such as:

(10) What is the number of planets in the solar system?

From external questions such as:

(11) Do numbers exist?

Carnap takes an internal question such as (10) as occurring from within an antecedently endorsed theoretical or linguistic framework, for instance the one of arithmetic. From within that framework, we can then ask questions as to the number of planets, *i.e.*, questions that make use of the framework in some way and that we answer by conducting empirical research. An external question such as (11), however, does not pertain to something we can ask from within the framework. Instead, it concerns the framework itself. That is, it asks whether the elements of the framework themselves, *e.g.*, numbers, exist. For Carnap, there is no "cognitive" answer to that kind of question (Carnap 1956: 207). On the contrary, the endorsement of a framework does not amount to an endorsement of a genuine ontology of numbers, but, say, to the endorsement of a set of "rules for forming statements and for testing, accepting, or rejecting them" (Carnap 1956: 208). One can either accept or reject the framework, but from Carnap's perspective it does not make sense to inquire about the reality of the framework's elements. Only internal questions are substantial.

We use the distinction between internal and external questions in a way that strongly opposes the idea that, with respect to the theory of intentionality, external questions are non-substantial. Indeed, one way to strictly apply Carnap's distinction to the problem of intentionality would be to claim something along the following lines: We do not have a grasp on the notion of aboutness from the outside of a specific theory of intentionality, *e.g.*, the CM. As such, we can only assess a certain theory of intentionality from an internal point of view. However, this is simply not true. Intentional experiences are not like numbers. Whereas we do not have a first-person access to numbers, we do have a first-person access to intentional experiences as we undergo them and as such have a first-person, pre-theoretical access to intentional experiences and to at least some of their features, *e.g.*, aboutness.

Of course, we do not necessarily have access to the *nature* of aboutness itself, which is something that might actually require a theoretical enterprise such as the one we are currently undertaking. Nonetheless, we do have access to aboutness as it is manifested in intentional experiences. As such, we do possess an external footing from which to assess theories of intentionality. An external question is then not a purely pragmatic one that concerns the way we decide to talk but, rather, a substantial one that concerns which theory of intentionality fits best the way intentional experiences are given to us through our first-person perspective.

A second remark about the above two kinds of questions concerns the relative degree of independence they enjoy. Indeed, the CM could be internally consistent while still failing to provide an adequate account of non-existence as evaluated from an external footing. As such, these two questions can be treated separately. We accordingly treat the internal question in the next section and the external one in section 5.

4. The Internal Challenge to the CM

The internal challenge to the CM can be summarized as follows. According to the CM, intentional experiences that are about things that do not exist have a content but no intentional object; they are about something that does not exist in the sense of having a content with some semantic properties that are independent from intentional experiences being successfully referential. In claiming this, however, the CM utilizes a substantial implicit premise, namely that it is possible to specify the semantic properties of a representational content that is about something that does not exist without making use of non-existent intentional objects. In other words, that it is possible for an intentional experience to be about something in the first sense of "about" as

distinguished above without thereby also being about something in the second sense of the term. How is that supposed to work?

Let us look at the following example:

(12) Sam judges that Pegasus is a horse.

As Pegasus does not exist, Sam's judgment is false. Nevertheless, Sam's judgment is about something—Pegasus, at least. According to the CM, this fact is supposed to be captured by the claim that the content of Sam's judgment, *Pegasus is a horse*, can be semantically evaluated despite not being referentially successful. However, what are the truth-conditions of the content *Pegasus is a horse*? Naively, one might think that they are the conditions in which there is a particular individual, Pegasus, who happens to be a horse. But how can we even start to specify such conditions if there is no horse, either existent or non-existent, which is identical to Pegasus? The CM must be able to provide an answer to that question.

The dialectic is then as follows. The CM might have a legitimate claim to the dissolution of the problem of non-existence as faced by the IOM, but it is confronted with a problem of non-existence of its own. Namely, how can we determine the truth-conditions of a representational content that is about something that does not exist in a way that does not appeal to non-existent intentional objects?

The consensus is that the CM can answer that challenge, although there is no established consensus about the particular semantic theory that it should endorse to do so. In the rest of this section, we identify some pros and cons for few options. However, we start with some general remarks on the notion of truth-conditions.

4.1 General Remarks on Truth-Conditions

We appeal to the traditional idea that an account of a formal language's semantic properties is the correct way to capture the notion of truth-conditions. For the time being, we also assume that the required formal language is either first-order classical logic or any of its extensions. Moreover, we assume familiarity with such a family of languages and merely recall one main explicit assumption that was already discussed in the previous chapter, namely that for any such language L there cannot be such a thing as an undefined individual constant. That is, for any such language L, there cannot be an element of the set of its individual constants {a, b, c, …} that is not assigned exactly one element of the domain of L.

Within first-order classical logic, we can distinguish between two distinct syntactic operations that give rise to a truth-evaluable sentence. The first is the *concatenation* of a predicative expression with an individual constant. This gives rise to so-called simple sentences of the kind: $F(a)$. Such sentences possess so-called *singular* truth-conditions; in other words, they are true if, and only if, a satisfies the predicate "being F." If a is not F, then the sentence is false. If a happens to be undefined, it is simply meaningless.

The second operation is the *prefixation* of a predicative expression by means of a quantifier, either universal or existential. Let us begin with the first case and assume that we are concerned with an unrestricted use of the universal quantifier. The prefixation of a predicative expression such as "F(x)" by means of the universal quantifier "∀" gives rise to a sentence of the form "∀(x)F(x)," to be read "Everything is F." This sentence is true if, and only if, any element of D satisfies the predicate "being F." If, on the other hand, such a predicative expression is prefixed by the existential quantifier "∃," this gives rise to a sentence of the form "∃(x)F(x)," to be read "Some x are F." This sentence is true if, and only

if, at least some elements of the domain are F or, alternatively, if the sub-set of D {x|Fx} is non-empty.

Importantly, the use of quantifiers can be combined to obtain more finely grained quantified sentences, that is, sentences that do not only tell us whether all or some x are F but actually how many, if any, x are F. One type of such a combination that plays a substantial role in the present discussion is the combination of the two quantifiers and the identity predicate "…=…" to express *uniqueness*, that is, the claim that at most one element of D satisfies a certain predicate, either simple or complex.

Such a claim of uniqueness can be expressed as follows:

(13) $\exists$(x) F(x) & $\forall$(y)F(y)—>x=y.

The first conjunct is true just in case {x|Fx} is non-empty while the second conjunct reveals that anything that is F is going to be numerically self-identical. The full conjunction hence indicates that the set *{*x|Fx*}* contains exactly one element and is true if, and only if, this is the case. Let us call such a sentence a *particularized* existential sentence.

Particularized existential sentences are to be distinguished from existentially quantified and universally quantified sentence even though all quantified sentences—universal, existential, and particularized—have the same kind of truth-conditions, namely so-called *general* truth-conditions. The idea behind this label is that these three kinds of sentences are true just in case some elements of D—either all, some, or a unique element—are F. However, for any instance of such a sentence type to be true, it is not required that one element of D in particular be F. That is, the very *identity* of the witness, or witnesses, does not matter; what matters is merely that some element of D—either all, some, or exactly one—satisfies some specific condition, namely F.

The difference between singular and general truth-conditions can be illustrated as follows. Assume that qualitative and numerical identity are distinct notions, that is, some things can be qualitatively identical without thereby being numerically identical. Now take the proposition that *a* is F. This proposition is true just in case a specific particular, *a*, is F. At some point, however, an omnipotent being destroys *a* and replaces it instantly by the qualitatively identical *b*. Can *b* make the proposition *a* is F true? No, granted that *a* is F has singular truth-conditions, nothing but *a* can make *"a* is *F"* true. Numerical identity, and not merely qualitative identity, matters. However, what about the particularized sentence (13)? This sentence would actually be just as true before and after switching from *a* to *b*. What matters here is mere qualitative identity.

Now that these preliminary distinctions have been made, we can reformulate the challenge the CM must face as follows. A proposition of the kind "Pegasus is a horse" intuitively seems to have singular truth-conditions as it seems to be saying something about Pegasus himself. However, as Pegasus does not exist, the sentence "Pegasus is a horse" cannot express the singular proposition "Being a horse (Pegasus)." How can the CM then assign some adequate truth-conditions to such a proposition?

Before seeing this, let us make a further clarification. We have now entered rather muddled territory. Our target is a theory concerning the content of intentional experiences. The contents of these experiences have a certain psychological reality. We are speaking about these intentional experiences by means of natural language, *e.g.*, English. Moreover, we have now introduced elements of a formal language, namely first-order classical logic, with the aim of more precisely capturing the logical form of the propositions expressed by the natural language sentences used to describe intentional experiences. One might legitimately wonder how all of these things mesh together.

The relation between these three things is complex and we cannot pretend to address the issue here. The best we can do is to clarify how we intend to have these three things working together, namely (i) the reality of intentionality as a psychological phenomenon, (ii) the assignment of propositional contents to intentional experiences, and (iii) the interpretation of these propositional contents in a formal language.

We conceive of these relations as follows. Intentional experiences are analyzed by the CM as non-intentional relations to representational contents. These contents might not necessarily be identified with propositions but it is at least helpful to treat them as such, in particular if truth-evaluability is the property of intentional contents we are tracking, as the framework of propositions is tailor-made to track truth. This is how we conceive of the relation between (i) and (ii).

Moreover, the right tool for accessing the nature of propositional contents is formalization. That is, we assume a standard distinction between grammatical and logical form. Grammatical form is the form of a sentence as parsed by the grammar of a natural language, whereas logical form is the form of a sentence as interpreted within a formal language. As already discussed in Chapter 3, these two things can sometimes come apart. The proposition expressed by the natural language sentence and, hence, its truth-conditions are revealed by the analysis of its logical form. This is how we conceive of the relation between (ii) and (iii).

We thus suppose a kind of two-step translation from the psychological reality of intentionality to natural language descriptions, and from natural language descriptions to formalizations. What is the nature of these translations? We remain agnostic: these translations either faithfully reflect the nature of intentionality or constitute a useful idealization. In either case, they provide us with a better understanding of some facets of intentionality. This is how we understand the relation between (i) and (iii).

To summarize, at this point we can understand the question that the CM faces as follows. The CM must be able to provide a specification of the truth-conditions of the representational content of the intentional experiences that it describes by means of natural language. Such a specification requires interpreting the sentences of natural language used to describe intentional experiences into a formal language. This interpretation yields sentences in a formal language that are taken to represent the form of the propositions expressed by the natural language sentences. The challenge then faced by CM proponents is to elaborate how they aim to account for the determination of the truth-conditions of representational contents that are about things that do not exist.

4.2 The Semantics of the CM

The CM must be able to determine the truth-conditions of representational contents that are about things that do not exist without appealing to the non-existent intentional objects. It has three main options for doing this:

(i) Endorse classical logic and attribute general truth-conditions to such contents;

(ii) Endorse classical logic and attribute singular truth-conditions to such contents; or

(iii) Endorse a non-classical logic and attribute non-classical logic truth-conditions to such contents.

We discuss the main tenets of these three options, thereby mentioning some of their main pros and cons.

4.2.1 Descriptivism

Option (i) amounts to endorsing some kind of descriptivist semantics, at least for singular terms that aim to refer to things that do not exist. In other words, it entails endorsing a semantics for such singular terms that treat their semantic contribution as equivalent to that of a particularized quantified sentence. Accordingly, propositions expressed by means of such singular terms turn out to be quantified ones.

The origins of descriptivism can be traced back to the doctrine's father, Bertrand Russell, who is well-known for having conceived his theory of definite descriptions as a barrier against the postulation of non-existent entities. Russell's theory can be schematically presented as follows. First, it carves a distinction between grammatical and logical form. Second, it argues that the need to postulate non-existent entities results from blindly following the grammatical form of sentences that are about things that do not exist when interpreting them in a formal framework. Third, it argues that once the correct logical form of the propositions expressed by these sentences is uncovered, the need to postulate non-existent entities vanishes. We now examine in further detail how this program is carried out.

We are already familiar with the distinction between grammatical and logical form from earlier parts of this work, but let us recapitulate. First, grammatical form is a sentence's form or structure as uncovered by natural language grammar. In the sentence "Jones met a man," the natural grammatical structure comprises a verb ("to meet") and two noun-phrases ("Jones" and "a man"). In contrast, logical form concerns the nature of the propositions expressed by such a sentence. In accordance with the account of truth-conditions introduced earlier, this logical form can be either one of a simple sentence that results from the concatenation of an individual constant or one of a quantified sentence.

Grammatical form can sometimes be misleading. Russell belabors this point first with an analysis of *indefinite* descriptions, that is, expressions such as "a man." The sentence "Jones met a man" can be

used as an example. This sentence's surface grammar tells us that the sentence is formed by means of a verb and two noun-phrases. According to Russell, however, some people have been misled into believing that the logical form of the proposition expressed by this sentence mirrors its grammatical structure and thus amounts to something like: Meet (Jones, a man).

If this were correct, however, we would be faced with a puzzle. In this analysis, "a man" is taken to be an individual constant and individual constants have first-order elements of the domain as semantic values. If this analysis were correct, we would thus have to find an element of the domain that serves as the semantic value of "a man." It hence seems that if we are willing to treat this sentence as having such a logical form, we must postulate a specific semantic value for "a man," namely an indeterminate man who adds to the numbers of all of the men there are. However, according to Russell there is no need to postulate such a weird kind of entities as the content of an assertion of the sentence "Jones met a man" is not equivalent to the content of an assertion that Jones met a particular man. As he argues,

> The question of 'unreality', which confronts us at this point, is a very important one. *Misled by grammar, the great majority of those logicians who have dealt with this question have dealt with it on mistaken lines. They have regarded grammatical form as a surer guide in analysis than, in fact, it is.* And they have not known what differences in grammatical form are important. 'I met Jones' and 'I met a man' would count traditionally as propositions of the same form, but in actual fact they are of quite different forms: the first names an actual person, Jones; while the second involves a propositional function, and becomes, when made explicit: "The function 'I met x and x is human' is sometimes true"… This proposition is obviously not of the form 'I met x'. (Russell 1919: 68–emphasis added)

Russell's diagnosis is the following: We have become trapped into having to answer odd questions such as "Which first-order element of the domain is the semantic value of 'a man'?," due to being misled by the surface grammar of natural language sentences. That is, we treated as individual constants expressions that actually belong to another logical category. Once we realize this, assures Russell, we can simply dissolve the original problem into the correct logical analysis. Accordingly, the real logical form of the proposition expressed by a sentence such as "Jones met a man" is something like the following:

(14) $\exists$(x) Man(x) & Meet (Jones, x).

This analysis does not commit us to any element of the domain being an indeterminate man. It instead commits to the predicate "Being a man" and can be taken as the content of an assertion of the sentence that Jones met a man. This proposition is then true if Jones indeed met a man, whoever that man is. But at no point in the process do we need to appeal to an indeterminate man as the semantic value of the expression "a man."

However, Russell did not limit his approach to indefinite descriptions. He successively applied it to all natural language singular terms, with different exceptions at different times (which we ignore here on grounds of convenience). His subsequent target was the analysis of *definite* descriptions, such as "the golden mountain." According to Russell, the formal interpretation of natural sentences formed out of such definite descriptions also creates some difficulties. If treated as an individual constant, a definite description such as "the golden mountain" requires that we attribute it a first-order semantic value. However, as no such thing as a golden mountain exists, we would be forced to accept non-existent things as the semantic value of definite descriptions that are

about things that do not exist. Russell is again not satisfied with such an answer, as he makes it clear in the following memorable passage:

> For want of the apparatus of propositional functions, many logicians have been driven to the conclusion that there are unreal objects. It is argued, *e.g.*, by Meinong, that we can speak about 'the golden mountain', 'the round square', and so on; we can make true propositions of which these are the subjects; hence they must have some kind of logical being, since otherwise the propositions in which they occur would be meaningless. In such theories, it seems to me, there is a failure of that feeling for reality which ought to be preserved even in the most abstract studies. Logic, I should maintain, must no more admit a unicorn than zoology can; for logic is concerned with the real world just as truly as zoology is, though with its more abstract and general features… The sense of reality is vital in logic, and whoever juggles with pretending that Hamlet has another kind of reality is doing a disservice to thought. A robust sense of reality is very necessary in framing a correct analysis of propositions about unicorns, golden mountains, round squares, and other such pseudo-objects. (Russell 1919: 169)

According to Russell, such a robust sense of reality can be preserved by extending his analysis of descriptions from indefinite to definite ones. That is, by considering that the semantic contribution of a definite description is not equivalent to the semantic contribution of an individual constant but, rather, of a quantified expression. However, this quantified expression cannot be of exactly the same kind as in the case of indefinites, on pain of blurring the semantic distinction of uniqueness implied by the passage from the indefinite article "a" to the definite "the." In other words, Russell asserts that the quantified expression must be a considerably more complex one that results, as we have seen above, from the combination of two quantifiers and the

identity sign in order to be able express the uniqueness feature of definite descriptions. This can be achieved by interpreting a natural language sentence such as "The golden mountain is high" as follows:

(15) $\exists$(x) Mountain (x) & Golden (x) & High (x) & $\forall$(y)
　　　Mountain (y) & Golden (y) & High (y) —> x=y.

The left-hand side of the principal disjunct is formed of the existentially quantified sentence "$\exists$(x) Mountain (x) & Golden (x) & High (x)." It can be taken to mean that at least one thing is a golden high mountain and is true just in case there is indeed in the domain something that does possess these properties. The right-hand side, "$\forall$(y) Mountain (y) & Golden (y) & High (y) —> x=y," is the conditional statement that if anything is a high golden mountain, then that thing is identical with itself. The conjunction of these two things says that there is at most one thing that is a high golden mountain. As no such thing exists, however, the proposition is false.

　　　The general idea here is exactly the same as the one at play in the analysis of indefinite descriptions. Treating definite descriptions as quantified sentences scratches off any problematic reference to first-order elements of the domain. As such, a sentence such "The golden mountain is high" does not express a proposition in which "The golden mountain" serves as a singular term that picks out a particular element of the domain. Rather, it expresses a quantified proposition constituted only of predicates, such as "being golden," "being a mountain,", "being high", and "…=…", as well as quantifiers. As such, it is not true just in a case a particular element of the domain is a high golden mountain but, rather, just in case at most one thing in the domain, whichever it is, satisfies these predicates. As Russell puts it in the *Philosophy of Logical Atomism*:

It is of utmost importance to realize that 'the so-and-so' does not occur in the analysis of propositions in whose verbal expression it occurs, that when I say 'The author of Waverley is human,' 'the author of Waverley' is not the subject of that proposition, in the sort of way Scott would if I said 'Scott is human,' using 'Scott' as a name. I cannot emphasize sufficiently how important this point is, and how much error you get in metaphysics if you do not realize that when I say 'The author of *Waverley* is human' that is not a proposition of the same form as 'Scott is human'. It does not contain a constituent 'the author of *Waverley*'. (2010: 89)

It is noticeable in this quotation that Russell makes a distinction between definite descriptions (such as "the author of *Waverley*") and names (such as "Scott"), to be understood here as the natural language equivalents of individual constants. As he later reflected on the matter, however, Russell came to realize that it is debatable whether natural language proper names really correspond to logical individual constants.

Russell thought two main issues seem to argue against such an assimilation of natural language proper. First, the problem of empty proper names makes the semantic contribution of natural language proper names as problematic as that of empty definite descriptions. Second, the problem of the cognitive value of proper names, famously rediscovered by Frege (Frege 1997a), makes it dubious that names can be semantically treated as merely contributing their reference to truth-conditions. Russell's strategy here was to kill two birds with one stone by arguing that both problems could be economically solved by means of treating all natural language proper names as abbreviated definite descriptions. We need neither to introduce non-existent objects nor accept the existence of Fregean senses; we instead merely have to treat proper names as abbreviated definite descriptions.

The general idea is as follows. For each name, we can come up with a conjunction of predicates that uniquely characterizes whatever the

proper name aims to refer to. We then form a uniquely restricted quantified sentence with that conjunction of predicates and obtain what we want: a definite description that the proper name can be taken to *abbreviate*. Take the example of the proper name "Scott." We can track down properties that uniquely characterize the referent of that name, such as "being the author of Waverley." The same is true for "Pegasus", *i.e.*, "being a horse," "having wings," and "being famous." The corresponding definite descriptions are then constructed out of these predicates and quantifiers. Consequently, natural language sentences such as:

(16) Scott is human;
(17) Pegasus eats;
(18) Pegasus does not exist.

Are taken to respectively express the following propositions:

(19) ∃(x) Author of Waverley (x) and Human (x) &∀(y) Authors
       of Waverley (y) & Human (y) —> x=y.
(20) ∃(x) Horse (x) & Flying (x) & Eating (x) & ∀(y) Horse (y) &
       Flying (y) & Eating (y) —> x=y;
(21) ¬ ∃(x) Horse (x) & Flying (x) & ∀(y) Horse (y) & Flying (y)
       —> x =y.

According to Russell, such a treatment of proper names offers a solution to the two problems alluded to above. First, we do not need to postulate exotic referents for empty names as propositions such as (19)–(20) are constituted only of concepts. Second, the cognitive difference between co-referential sentences such as "Hespherus = Hespherus" and "Hespherus=Phosphorus" can be accounted for in terms of the fact

that the two proper names are given different definite descriptions that then turn out to be co-extensional.

With these standard elements of descriptivism in place, we can formulate a first way in which the CM may be able to account for the truth-conditions of the representational contents of intentional experiences that are about things that do not exist: namely, for any representational content of an experience that is about something that does not exist, this content is accounted for along descriptivist lines. For instance, if John judges that Pegasus eats, the content of this experience, namely *Pegasus eats*, can be analyzed along the lines proposed above. As such, descriptivism would allow the CM to saddle such experiences with contents without thereby having to appeal to anything such as Meinongianism. As John McDowell summarizes the issue:

> The point of the theory of descriptions is exactly to avoid an apparent need for nonexistent real objects as relata for intellectual acts. Where a relational conception of intellectual acts would require them to stand in relations to possibly non-existent objects, Russell instead takes their content to include *specifications* of objects. If no objects answer to the specifications, that does not threaten the contentfulness of the acts. (McDowell 2009: 23)

How successful is this first option? In such a simple version, descriptivism is threatened by a battery of arguments that are much discussed in contemporary semantics. These arguments generally target a descriptivist treatment of referential proper names, although at least some of them can also be extended to empty names. For instance, in his influential *Naming and Necessity*, Saul Kripke is well-known for having proposed three powerful arguments against descriptivism: the so-called modal, semantic, and epistemic arguments (Kripke 1980). We focus here on the the semantic argument.

The semantic argument focuses on the descriptivist claim that for each proper name *N*, *N* refers, as a matter of its semantics, to whatever is singled out by its descriptive content. However, Kripke contests that this is the case by means of the following example:

> Suppose that Gödel was not in fact the author of [Gödel's theorem]. A man named "Schmidt," whose body was found in Vienna under mysterious circumstances many years ago, actually did the work in question. His friend Gödel somehow got hold of the manuscript and it was thereafter attributed to Gödel. On the view in question, then, when our ordinary man uses the name "Gödel," he really means to refer to Schmidt, because Schmidt is the unique person satisfying the description, "the man who discovered the incompleteness of arithmetic"... So, since the man who discovered the incompleteness of arithmetic is in fact Schmidt, we, when we talk about "Gödel," are in fact always referring to Schmidt. But it seems to me that we are not. (Kripke 1980: 83-84)

The gist of this example can be summarized as follows. If descriptivism is true, then a proper name *N* refers to what is singled out by its corresponding definite description. Moreover, if it is true that the referent is thus singled out, then it follows that what a competent speaker refers to by means of a use of a proper name *N* must be a function of the beliefs that she holds about the referent of *N, i.e.,* these predicates the speaker believes to be true of whoever, or whatever, uniquely satisfies them. However, the above Gödel example demonstrates that this cannot be true. Indeed, "Gödel" does not refer to whoever satisfies some description we standardly attach to "Gödel," such as "being the man who discovered the incompleteness of arithmetic." If this were the case, then in a counterfactual situation in which it would be false that Gödel discovered the incompleteness of arithmetic, "Gödel" would refer to the person who really discovered the incompleteness of

arithmetic, *e.g.*, Schmidt. However—and here is Kripke's central claim—
we keep referring to Gödel by means of "Gödel" even if the
descriptions we attach to "Gödel" are false. As such, the reference of
"Gödel" cannot be determined by some definite description.

This argument is interesting in that it also seems to be
straightforwardly applicable to non-referential proper names. Imagine
that we coin a non-referential proper name *N* to talk about a non-
existent individual who we describe by means of a definite description
"The *F*." If descriptivism is true, then "The *F*" determines the reference
of *N*. Unbeknownst to us, however, there might be someone who exists
and satisfies the description "The *F*"; let us call this person *a*. However,
N clearly does not refer to *a*; it instead aims at referring to that thing we
know does not exist.

The goal of this argument is not to prove descriptivism wrong; it
is rather to show that if the CM is being willing to endorse a descriptivist
semantics, this comes at a certain theoretical cost, namely the need to
find rejoinders to this and related arguments. That being said,
proponents of descriptivism must beware of the following difficulty: the
more complex descriptivism becomes, the less plausible it is that the
common usage of singular terms can be explained in descriptivist terms.
This is because it seems unlikely that common speakers will routinely
treat singular terms as equivalents to philosophically complex definite
descriptions.

Moreover, one might have the intuition that whenever one is
thinking when, say, one is thinking about Pegasus, the content of one's
intentional experiences goes beyond "specifications" of Pegasus, to use
McDowell's phrase. That is, whenever one is thinking about Pegasus, one
tokens a kind of thought that seems similar to a thought about an
existent object such as, say, Obama. The CM, however, must either treat
both kinds of thought as equally having a content that primarily amounts
to specifications of particulars or bite the bullet that thoughts about, *e.g.*,

Pegasus are of a different kind than thoughts about, *e.g.*, Obama. One might argue that both options have unpleasant consequences.

4.2.2 *Non-Descriptivism*

On the basis of the above remarks on descriptivism, one might be willing to explore the traditional alternative to descriptivism: non-descriptivism. However, non-descriptivism comes in two different flavors. The first one, which we refer to as "Millianism"—after John Stuart Mill—asserts that a proper name's semantic content is exhausted by its referent. The second option, which we call "non-Millian non-descriptivism" (NMND), contends that a proper name's semantic content is non-descriptive but does regard it as being exhausted by its referent. Let us briefly discuss each option.

The Millian option probably constitutes the less natural ally to the CM. Indeed, if some intentional experiences are about things that do not exist and to which we aim to refer by means of proper names, then Millianism will not be able to account for their representational content. However, in most of such cases, Millians are willing to take a hard line and claim that a great number of what we take to be intentional experiences that are about things that do not exist are in fact about things that do exist. First and foremost, Millians standardly recognize the existence of fictional characters in order to serve as the semantic value of fictional proper names (Van Inwagen 1977; Salmon 1998; Thomasson 1999; Voltolini 2006). We cannot evaluate these arguments here, in particular because most of them are independent of considerations that pertain to intentionality; nonetheless, we point to two important issues. First, this leaves open the exact treatment that Millianism should give of proper names for which even Millians themselves do not want to concede existent referents. Most importantly, someone ready to embrace a conjunction of the CM and Millianism would have to live with an

important and daunting drawback, namely the simple observation—that non-descriptivist would certainly dismiss as philosophically naive—that many people would simply consider the claim that fictional characters and the like exist as ontologically unacceptable and phenomenologically inadequate. Indeed, as Jody Azzouni remarks:

> Imagine a child pointing at a cartoon figure while watching television. "That's Santa Claus," the child shouts gleefully. One day in the near future, she will revise her ontological commitments: She will realize (to her dismay) that there is no Santa Claus. We adults –most of adults, anyway– already know that there is no Santa Claus. What, then, was she referring to when she said, "That's Santa Claus"? It seems that the answer should be this: She was referring to Santa Claus. And it should be added that Santa Claus is nothing at all–that is, it should be added that Santa Claus doesn't exist. (Azzouni 2012: 253-254)

On the other hand, NMND is often formulated by drawing a Fregean distinction between the sense and reference of a proper name. Within that framework, however, senses are deemed to be non-descriptive. In other words, senses reflect the cognitive contribution of a proper name, in addition to the semantic contribution made by its referent—although this cognitive contribution is not to be identified with a definite description that determines the referent. In recent years, an increasingly influential idea is that such non-descriptive senses can be identified with so-called "mental files" (Recanati 1993, 2010, 2012).

The key idea behind the notion of mental file is that our cognitive life is structured by the opening and updating of mental files, either temporary or more permanent, that pertain to individuals about which we entertain intentional experiences. The files may be descriptive or non-descriptive, with the distinction cashed in the way the different kinds of files treat information about their referents. In the case of a descriptive

file, the referent is always determined *satisfactionally*, *i.e.*, as a function of the information contained in that file. As such the file's referent can change as the information in the file is updated. For instance, our file labeled "The president of the USA" now refers to Obama but will refer to Trump as soon as he, alas, takes office. On the other hand, non-descriptive files contain information about their referents but this information does not serve to determine these referents. The reference of non-descriptive files is not determined satisfactionally but *relationally*, by what Recanati calls "epistemically rewarding" acquaintance relations (Recanati 2010: 156), such as perception, testimony, and the like.

Of course, in the case of intentional experiences that are about things that do not exist, there cannot be such relational determination of reference. However, the proponents of NMND assert that this does not mean that such experiences must be saddled with descriptive contents; they instead assume that there is a middle ground between descriptive contents and fully-fledged non-descriptive referential contents. These contents can be characterized as the opening of non-descriptive files that fail to have a referent. Such openings are somewhat degenerative openings of genuine non-descriptive mental files, and are parasitic on them; nonetheless, they are fully non-descriptive.

This idea, however, is not too promising for the CM. First, the CM needs representational contents with truth-conditions in order to account for aboutness, and non-descriptive files that lack a referent would not be able to determine such truth-conditions. These files might play a certain cognitive role, but they cannot play a semantic role. Second, it is hard to make phenomenological sense of this notion of mental files. The talk of opening mental files might make sense at a sub-

conscious level, but it seems out of place in a description of intentional experiences.[20]

In sum, then, both descriptivism and non-descriptivism do face some problems and there is no easy, straightforward, road for the CM to take. As such, then, the IOM is not the only theory that needs to pay a theoretical price to be able to account for non-existence. The CM must as well pay some theoretical price, though in a different, semantic, currency by embracing either some descriptivist or non-descriptivist semantics, each with some respective problems and challenges.

According to some philosophers, however, we might be able to avoid an oscillation between descriptivism and non-descriptivism if we recognize the fact that at the core of both positions lie an assumption about the kind of formal framework in which semantic inquiries should be conducted. Once this assumption is bracketed, however, these philosophers claim that we could come up with an account of the semantic of proper names that would avoid much of the difficulties encountered by the two standard positions. Moreover, such a rejection of classical logic would not necessarily amount to an endorsement of Meinongianism. In the next section, we quickly present the main tenets of an account willing to go down this alternative road: namely, the one proposed by Mark Sainsbury in his book *Reference without Referents* (Sainsbury 2005).

---

[20] Of course, we adapt here the framework of mental files to intentional experiences but Recanati conceives of intentionality as being essentially a matter of sub-conscious states. As such, then, the objection is not targeted at Recanati's himself but, rather, at the application of the idea of mental files to a conception of intentionality as being essentially conscious.

4.2.3 The Temptation of Non-Standard Semantics

As already mentioned earlier, classical logic treats it as axiomatic that there are only two ways to form a meaningful, *i.e.*, truth-evaluable, sentence: either by concatenating a well-defined individual constant with a predicate or by prefixing a predicate by a quantifier. Everything else, *e.g.,* use of an empty individual constant or a mere open sentence, has no meaning—or in the latter case, no meaning on its own (*cf.* Sainsbury 2005: 45).

Aside from the purely formal reasons that exist for supporting these axioms, their endorsement is also quite intuitive. As Russell puts it: "Logic is concerned with the real world just as truly zoology is" (Russell 1919: 169). It hence seems natural that logical sentences are formed in such a manner that they are anchored in reality right from the start, by means of either individual constants that stand for elements in the domain or quantifiers that range over things that exist.

Nonetheless, this claim to naturalness can be challenged. A classical objection is that logic should be a completely *a priori* discipline and that the claim that individual constants must all be well-defined is an *a posteriori* one. Indeed, classical logic presupposes that the domain of quantification is not empty, which is something that logic alone cannot account for. This can be argued as follows.

In classical logic, it is a matter of logical truth that for any individual constant *a*, "$\exists(x)\ x = a$" is true. Granted that proper names are the equivalent of individual constants, "$\exists(x)\ x = \text{Obama}$" is hence true on purely logical grounds. Since this sentence is logically true, then any substitution of "Obama" by a syntactically similar term should yield a truth. However, "$\exists(x)\ x = \text{Pegasus}$" is false, as Pegasus does not exist. As a result, the truth of "$\exists(x)\ x = \text{Obama}$" is a logical one only under that assumption that Obama exists. However, this is not a logical claim as one cannot know on purely logical grounds that Obama exists.

As Karel Lambert puts it:

In general, no singular term that falsifies the condition

　　s exists,

can be allowed to replace the singular term ["Obama"] in [the above sentence]. So a logic containing that statement among its logical truth is not free of existence assumptions with respect to its singular terms because, again, there are statements in it that will not be logically true unless it is true that **s** exists for any singular term **s**. (Lambert 2003: 125)

However, if logic is supposed to be a purely *a priori* discipline, it cannot be taken as axiomatic that individual constants must, as a matter of logic, be assigned existent referents. As Sainsbury puts it:

Applying classical logic to language would require a prior segregation of its proper names into those which have bearers and those which do not; only the former should be allowed to replace individual constants in the classical formalism. Effecting this segregation does not seem to be part of the job of a logician (traditionally conceived): logic is supposed to be *a priori,* and not to involve the kind of astronomical or literary knowledge require to determine in which category a name like 'Vulcan', Homer', or 'Patanjali', should be placed. (Sainsbury 2005: 65)

This objection is not tied to the problem of either non-existence or empty terms; it is an objection about the foundations of logic itself. If correct, the assumption that classical logic constitutes a privileged formal language for semantic inquiries appears to be ungrounded.

Another objection more closely related to the issues addressed in this chapter that one could address to proponents of classical logic is

that the use to which we put logic is an inquiry into not only "the more abstract and general features" of reality (Russell 1919: 170) but also the *means* we use to speak about and represent reality. Moreover, these means have many uses, including the use of empty names. As such, our privileged formal language should not be one that right from the start ostracizes a non-negligible portion of natural language. In light of this remark, the fact that we get stuck with an oscillation between descriptivism and non-descriptivism might be viewed as a sign that this remark has not been taken seriously enough and that we are operating within the boundaries of a logic axiomatized too restrictedly.

Let us call any logic that rejects the axiomatization of classical logic non-classical. An important family of non-classical logics are the so-called "free logics," which are called thus as an abbreviation for the phrase "free of existence assumptions with respect to its terms, general and *singular*" (Lambert 2003: 123). In essence, these are logics in which the above-mentioned condition "**s** exists" is lifted for the formation of simple sentences. Within a free logic, the criterion for the formation of a meaningful simple sentence is hence a mere syntactical one, free of any assumption about what the domain contains. In other words, once a set of predicates {F, G, H, …} and a set of individual constants {a, b, c, d, …} have been provided for a language *L*, one can meaningfully combine their elements into well-formed sentences independently of any assumption with respect to these sets' ontological commitments. [21]

A direct consequence of endorsing a logic that countenances such sentences as well-formed concerns quantifiers. This consequence, however, takes different forms depending on the kind of free logic one is willing to endorse. In Chapter 3, we saw how a Meinongian free logic must introduce unloaded quantifiers to quantify over its outer-domain of

---

[21] We ignore here the argument in favor of the claim that predicates should not also be considered as ontologically committed. See, for instance, Lambert 2003.

quantification; however, other kinds of free logics deal with quantification differently.

The kind of free logics we are concerned with in this chapter do not aim to quantify over non-existent objects, unlike a Meinongian free logic; as such they do not need to introduce unloaded quantifiers. Nevertheless, the rejection of the existence assumption placed on individual constants by classical logic forces such a kind of free logic to revise the standard quantifier rules of universal instantiation (UI) and EG.

The UI rule specifies that from any true universally quantified sentence of the form "$\forall(x)\ F(x)$," one can validly infer the truth of any sentence that results from substituting the variable "x" with an individual constant. On the other hand, EG is the rule that from the truth of a sentence of the form "F(a)," one can validly infer the truth of an existentially quantified sentence of the form $\exists(x)F(x)$. However, both of these rules must be modified in a free logic that maintains a classical reading of the quantifiers. Indeed, since empty individual constants are allowed, there is no guarantee that an occurrence of an individual constant will refer to something that falls within the scope of the quantifiers, classically conceived. Accordingly, these rules must be modified as follows:

> *Universal Instantiation*: From $\forall(x)$ and $\exists(x)\ x = t$ infer A(t/x) (where 'A(t/x)' is the formula that results from 'Ax' by replacing every occurrence of 'x' by 't').
>
> *Existential Generalization*: From A(t/x) and $\exists(x)\ x=t$ infer $\exists(x)$A(x).
>
> (Sainsbury 2005: 65)

As such, UI is modified as follows: One cannot infer that, *e.g.*, t is A from $\forall(x)$ A(x), but rather only from the conjunction of $\forall(x)$ A(x) and the claim that t exists, i.e., $\exists(x)\ x = t$. On the other hand, EG is modified as

follows: one cannot infer that $\exists(x)A(x)$ from $A(t)$ alone, but rather only from the conjunction of $A(t)$ and the claim that t exists, *i.e.*, $\exists(x)\ x = t$.

Once the idea of a free logic has been accepted, however, one is faced with an important question that concerns the possible truth-values of sentences now licitly formed by means of an empty name. Here a free logic can take one of the following three forms (Sainsbury 2005: 66):

- **(i) Negative free logic**: all simple sentences that result from the concatenation of an empty name and a predicate result in a falsehood;
- **(ii) Non-valent, or Fregean, free logic:** sentences that result from the concatenation of an empty name and a predicate are well-formed but truth-valueless;
- **(iii) Positive free logic:** at least some simple sentences that result from the concatenation of an empty name and a predicate result in a truth.

In this chapter, we are only concerned with (i) as this is the version that Sainsbury endorses to frame his theory. In passing, however, note that (iii) is not to be identified with a Meinongian free logic of the kind discussed in Chapter 3. Indeed, a Meinongian free logic is certainly positive as it recognizes that there can be true predications about things that do not exist. That being said, the endorsement of unloaded quantifiers is not a necessary feature of a positive free logic; the recognition that there can be true predications formed out of empty individual constants, however, does constitute such a condition.

In his book *Reference Without Referents*, which was inspired by pioneering work by Tyler Burge (Burge 1974), Mark Sainsbury has proposed a semantic framework based on negative free logic that he asserts,

> …is not Millian, for it allows that there are intelligible empty names, but it is also not descriptivist in that it denies that any names have a meaning that is properly captured by a definite description. (Sainsbury 2010: 39).

We present the great lines of his proposal below.

Sainsbury's fundamental claim is that all instances of proper names do not get assigned as their fundamental semantic value neither a referent nor a description but, rather, a non-descriptive reference condition of the form:

(20) $\forall$(x) "Obama" refers to x iff x = Obama.

Such a condition distinctively appeals to (i) signs, such as "Obama"; and (ii) the property of being identical with Obama, formally represented as "x=Obama." However, this condition is not descriptive, for the following reason. As Sainsbury puts it, the aim of such a condition is not "an attempt to analyze the meanings of individual words" (Sainsbury 2005: 74), which is a task that descriptivism avowedly pursues. That is, as a statement of the meaning of "Obama," a condition such as (20) is idle as the word "Obama" is present on both on left- and the right-hand sides, once mentioned and once used. In no sense this would suffice as an analysis of the meaning of the word "Obama."

According to Sainsbury, such axioms can nonetheless accomplish a great deal of work. For instance, they can assign a semantic contribution to empty names. That is, for any empty name, *e.g.*, "Pegasus," such a condition will turn out to be false as nothing corresponds to it. Nonetheless, "Pegasus" can be used to form simple sentences that will be semantically evaluable, which is impossible in a classical logic set-up.

As a result, for any predicate *F*, *"*Pegasus is F" will turn out false as a matter of the axiomatic of classical logic. Sentences that are formed out of empty names can nonetheless be true when embedded into an operator, either a truth-functional one such as "¬" or an intensional one such as "believes that." For instance, if "Pegasus exists" turns out to be false, which is fine, its negation, "¬(Pegasus exists)" will turn out to be true, which is even better. Similarly, if "Sherlock Holmes lives 221B Baker Street" is false, it could turn out to be true if embedded in an intensional operator such as "John believes that" or "In the Sherlock Holmes fiction."

Note also that contrary to descriptivist semantics, the semantic contribution of these referential axioms for empty names is singular. Indeed, since for any *x* there is at most one thing that can be numerically identical with it, if the referential axiom turns out to be true it is that thing, and not any other one, that makes it true.

Sainsbury's underlying idea is that such referential axioms capture the tacit knowledge of truth-conditions that comes with minimal semantic competency with a proper name, which is a general competency that one acquires when learning a natural language. A genuine understanding of instances of proper names comes later progressively as one becomes immersed within what Sainsbury calls a "proper name-using practice" (Sainsbury 2005: 106).

If Sainsbury is on the right track, then we would have means to assign truth-conditions to propositions expressed by means of empty names in such a way that, first, they would not require the assignment of a referent to be meaningful and, second, their meaning would not end up being equivalent to a definite description. As such, it seems that Sainsbury's position constitutes an interesting middle ground for the CM.

Nonetheless, Sainsbury's framework is not without problems. First, it is not obvious that all true sentences expressed by means of an empty proper name are sentences that result from the embedding of a

false sentence under an intensional operator. For instance, the meta-fictional sentence "Sherlock Holmes is more famous than Miss Marple" (Crane 2013: 81) is true and formed out of two fictional proper names. However, it is doubtful that there is an intensional operator under which it is embedded and that makes it true.

A second problem concerns non-propositional attitudes, such as "to think about" (Sainsbury 2010). Such verbs are transitive and take a direct grammatical object as a complement, as in: "I am thinking about Pegasus." Now, imagine that this is true, *i.e.*, that I am indeed thinking about Pegasus. According to Sainsbury's account, this cannot be literally true as Pegasus does not exist; however, this certainly is literally true. Sainsbury's account thus seems to be extensionally inadequate. Of course, one might protest that there is no such thing as a non-propositional attitude; we come back to this point in more detail in the next section.

5. Evaluating the CM

At the beginning of the preceding section, we made a distinction between two kinds of questions one can address to the CM: internal and external ones. The central internal question was whether the CM is able to determine the truth-conditions of the representational contents of intentional experiences that are about things that do not exist. As we have seen, this issue is complex and there is room for much philosophical debate. Nonetheless, proponents of the CM are confident that some adequate account of the semantics of representational contents that are about things exists. We mentioned in passing some main bullets they might have to bite along the way.

In this section we turn to external issues. In this regard it should be recalled that our distinction between internal and external questions was fundamentally different from Carnap's. Indeed, Carnap contended

that external question are non-substantial as he thought that there is no external footing from which to address substantial questions about frameworks. However, we assert that the case of intentionality is different: intentional experiences have peculiar features: they are phenomenological entities that have a psychological reality and that can be accessed from the first-person perspective; a good theory of intentionality should be accountable to these features.

We mention two central external objections below. The first concerns the existence of irreducibly non-propositional—or objectual—attitudes while the second concerns the notion of aboutness itself.

5.1 Non-Propositional Attitudes

The first objection that we consider is based on the existence of so-called non-propositional attitudes, that is, intentional attitudes that are *irreducibly* non-propositional. These attitudes are described by means of an intentional transitive verb, *i.e.*, an intentional verb that takes a direct, rather than indirect, object. However, the point is not strictly speaking linguistic but phenomenological. Namely, there is a real phenomenological distinction between non-propositional and propositional attitudes and this distinction is reflected at the grammatical level.

For instance, consider the following examples:

(21) John loves Mary.
(21) John is seeking his car.
(23) John contemplates the moon.
(24) John thinks about Obama.
(25) John notices Samantha.
(26) John admires Obama.
(27) John fears God.

All of these examples initially share the same grammatical form: they all feature an intentional transitive verb, *i.e.*, a verb that takes a direct object. Beyond their mere grammatical form, they also seem to share a phenomenological feature: they all are instances of intentional experiences directed upon particulars, without the presence of an instance of predication. That is, to love Mary is undergo an emotion directed at Mary herself and not necessarily to love that Mary possesses some feature. Similarly, for John to seek his car is not for John to seek that his car is somehow. It is simply to seek for a particular.

That being said, at the outset we have to distinguish between two issues, namely:

(i)   The notion that intentional experiences are always directed upon something under a certain aspect; and:

(ii) The notion that intentional experiences are always propositional.

Indeed, to defend that some intentional attitudes are irreducibly propositional does not entail endorsing the idea that objects of non-propositional attitudes are not being intentionally targeted under a certain aspect. For instance, to defend the claim that an attitude of the kind "John thinks about Obama" is irreducibly non-propositional does not entail that John is not thinking about Obama—the man himself, as could be pointed to from the point of view of nowhere, as Nagel would put it (Nagel 1986)—*as* Obama or, similarly, that when John notices Samantha, he notices her under some aspect: namely, that aspect of Samantha looking some way from John's perspective. Rather, what one defending such a view is committed to is the claim that there is a difference between, on the one hand, being intentionally directed at an object under a certain aspect and, secondly, being intentionally directed at something in a propositional way.

Where does the difference lie? Two complementary features are worth pointing to. First, as mentioned above, not all intentional experiences are predicative ones. That is, not all intentional experiences are experiences that have, as part of their nature, the application of a predicate to a particular. For instance, to undergo the intentional experience as of being presented with the tea cup being on the left seems to entail somehow that one attributes, consciously or sub-consciously, the property of being on the left to the tea cup. Accordingly, the content of that experience can be expressed by saying that one experiences *that* the tea cup is one the left.

However, it does not seem that all intentional experiences are thus predicative. John's intentional experience of loving Mary, for instance, is not predicative. Of course, John could find *that* Mary is lovable, which would be a clear instance of an intentional experience that entails a predicative part. However, to love Mary and to find Mary lovable are not the same kind of experience, as one may find lovable many persons one does not, in fact, love. Accordingly, there is a basic difference between being non-propositionally intentionally indirected upon something under some aspect and, second, being propositionally directed upon something.

A complementary point is that not all intentional experiences we undergo are experiences that have a semantic nature, *i.e.*, not all experiences are experiences that the world *is* or *should be* some way. Some experiences are simply appearances of things. Propositional attitudes can be semantically evaluated: judgments can be true or false, desires can be satisfied or non-satisfied, perceptions can be veridical or non-veridical. However, non-propositional attitudes are not the same. For instance, John's contemplating the moon is not semantically evaluable. It is just an experience as of John turning his gaze in some way to the moon.

On the basis of these observations, we may then say that there is a good *prima facie* case in favor of the existence of irreducibly non-

propositional attitudes. Michelle Montague aptly summarizes the matter as follows:

> Simply put, objectual attitudes resist a propositionalist analysis. Mary loves Nancy. She seeks the fountain of youth. She has you in mind. She contemplates the sky. And she wants Nancy's car. These intentional attitudes appear to be relations that hold simply between thinkers and non-propositional objects, rather than between thinkers and propositions, and we need to be given some reason to think otherwise. The burden of argument is plainly on the propositionalists. (Montague 2007: 507)

So far, the case in favor of the existence of non-propositional attitudes is only *prima facie*; unless one is given some reason to think otherwise, it just appears that there are non-propositional attitudes. However, the CM gives us an important reason to think otherwise. Indeed, the CM provides an analysis of intentionality that relies on the idea that for an intentional experience to be intentional is for it to have a certain representational content from which we can derive its property of being about something. The existence of non-propositional attitudes threatens this analysis. The problem that the existence of irreducibly non-propositional attitudes creates for the CM can be then summarized by means of the following two incompatible theses:

(A)   Some intentional attitudes are irreducibly non-propositional; and

(B)   The intentionality of an attitude is derived from the semantic properties of its propositional content.

As noted, (A) and (B) are incompatible. Indeed, if (A) is true, then there are some intentional attitudes that do not derive their intentionality from

the semantic properties of their propositional content. On the other hand, if (B) is true, then there cannot be irreducibly non-propositional intentional attitudes. The present objection to the CM argues that we should endorse (A) and reject (B).

However, the CM might try to resist this objection by rejecting (A) and the *prima facie* case in its favor provided above. Indeed, it might claim that even though it might *appear* that there are non propositional attitudes, these are in fact reducible to propositional ones. How could the CM argue in favor of that claim?

A CM proponent could first employ a syntactic strategy to pursue that aim. That is, she could try to defend the claim that what we take to be non-propositional *sentences* in fact possess hidden propositional clauses upon analysis. One tactic employed in the literature to achieve this consists of modifying a sentence that contains an intensional transitive by means of an adverb and then showing that this modification creates ambiguities that can be resolved only by uncovering covert propositional material. For instance, consider the following sentence:

(28) John wants your bike tomorrow.

This sentence is ambiguous: it can be read as meaning either that John *now* wants your bike for tomorrow or, alternatively, that John will *tomorrow* want your bike, without necessarily requiring your bike for tomorrow. In the first case, "tomorrow" modifies the content of John's wanting, that is, it is a specific wanting that John has—a wanting for tomorrow. In the second case, however, "tomorrow" is not used to modify the content of the wanting but, rather, the time at which it will occur in John's mind.

A proponent of the CM can then claim that in order to be able to account for this ambiguity, one needs to treat (28) as containing a hidden

propositional construction. In this regard, we can then disambiguate (28) as follows:

(29) John is such that he wants your bike for tomorrow.
(30) Tomorrow, John will be such that he will want your bike.

The CM proponent can then try to demonstrate that this strategy can be applied to all intentional transitive constructions, which would demonstrate that all non-propositional constructions are in fact covert propositional constructions. For instance, Richard Larson writes the following:

> In fact, a 'hidden clause' analysis of intensional transitive constructions is plausible in many cases, and has been urged by a number of researchers. Quite typically, intensional transitive constructions have a close paraphrase involving a clausal, or clausal-like construction. (Larson 2002 : 233)

However, the CM proponent must confront two problems with such a syntactic strategy. The first is simply that this constitutes a mere syntactic observation and that it is not clear exactly how it impinges on the phenomenological data described above. Indeed, does the fact that (28) is ambiguous really show that a simpler instance of intentionality, such as the one described by "John wants your bike," is not an instance of a non-propositional attitude? The proponent of the CM must here give us some reason to believe that this is the case. She cannot simply put forward the fact that (28) is ambiguous and requires a hidden propositional structure to account for this ambiguity as an argument in favor of the claim that there are no irreducible non-propositional properties.

The second problem is that this syntactic strategy in fact does not generalize. Indeed, as remarked by Montague, such an ambiguity test works well for verbs from the "want/needs" verb group, but it does not generalize smoothly beyond that group (Montague 2007: 508). For instance, it is indeed difficult to distinguish any ambiguity in the following:

(31) John loves Mary tomorrow.

One might imagine that one is anticipating John's first encounter with Mary and, aware of John's sensibility as well as Mary's natural charm, knows that John will instantly fall in love with Mary tomorrow. However, there is no way to read this sentence as saying that "tomorrow" somehow modifies John's loving, as was the case with (28) above.

Alternatively, one may point to an example such as:

(32) John loves Mary faithfully.

Here it seems quite clear that "faithfully" modifies John's loving of Mary, but it is once again hard to find any ambiguity in the sentence.

Montague suggests an interesting reason as to why the "wants/ needs" verb group is, unlike other groups of transitive verbs, subject to such ambiguities. Desire is commonly taken as being directed at a way the world can be, *i.e.*, as a paradigmatic case of a propositional attitude. The proximity that members of the "wants/needs" verb group have to desire might then explain why such constructions in fact contain hidden propositional clauses.

This first syntactic strategy does not seem to be very promising for the CM. How else can one try to defend the claim that non-propositional attitudes can be reduced to being propositional? We might try to saddle a proponent of the CM with the following strategy,

discussed by Montague 2005. First, we might say that for each non-propositional attitude, there is a true bi-conditional with a non-propositional attitude on one side and a propositional attitude on the other; for instance:

(13) John fears God *iff* John fears that God is omnipotent.

The problem with such a bi-conditional is that if the right-hand side certainly counts as a sufficient condition for John to fear God, it does not constitute a *necessary* condition. Indeed, John could well ignore that God is omnipotent or even be unfamiliar with the concept of omnipotence and yet still fear God. Of course, the point is not merely about (13), but rather about the way we can generalize over this example. Namely, it seems that if the content of the propositional attitude that stands on the right-hand side of the bi-conditional is too specific, we could always remark that such an attitude might not count as a necessary condition for the non-propositional one on the left-hand side.

However, we might retreat from the attribution of a determinate content and merely claim that for any non-propositional attitude of the form "*x φs y*" there is some propositional attitude with the form "For some *F*, *x φs* that *y* is *F*" that is true; moreover, we could also assert that the two propositions that would be instances of these two schemas could truly flank a bi-conditional. For instance, one may claim that if "John loves Mary" is true, then the following bi-conditional is true: "John loves Mary *iff* for some *F*, John loves that Mary is *F*." However, here it seems that we might be confronted by the converse of the above problem, namely that the truth of the propositional attitude that would flank the right-hand side of the bi-conditional would not count as being *sufficient* for the truth of the propositional attitude. That is, John might well love that, for some *F*, Mary is *F*, without this entailing that John loves Mary. If that is not the case, then CM proponents should be able

to tell us which *F* exactly are such that they are sufficient for the fact that if Mary has one of them, then John loves her. However, the prospects for the existence of such a property appear rather dim.

As such, we might also conclude that this second strategy is not more promising than the first and that, consequently, it does not seem that we have good reasons to reject (A). The existence of irreducible non-propositional attitudes thus does constitute a genuine problem for the CM.

5.2 Aboutness

We now turn to a second external objection to the CM. In essence, it amounts to the claim that the CM's analysis of aboutness is wrong-headed. In order to properly understand this objection, however, it is important to distinguish between two versions of it.

First, one might say that from the perspective of an IOM proponent, the analysis of aboutness suggested by the CM is insufficient. Indeed, for proponents of the IOM, aboutness is defined in terms of the notion of intentional object. By rejecting an irreducible appeal to that notion, the CM would—from the point of view of the IOM— simply be missing the nature of aboutness.

Of course, this first way to frame the objection is not very interesting. Indeed, as we saw earlier its antecedent, namely that the IOM provides us with a correct analysis of aboutness is rather problematic. Nonetheless, one might be willing to distinguish between two things: the *general intuition* that lies behind the analysis of aboutness provided by the IOM and the *specific analysis* of this intuition of the IOM.

What is this general intuition? To put it in Crane's terms, it is that "our mental life seems to involve the presence to the mind—or apparent presence to the mind—of things in the world" (Crane 2013: 4). A compelling reading of what Crane is saying is that for each intentional

experience, there is something such that it appears that we are intentionally directed upon that thing. However, the CM is bound to reject this intuition. As we saw above, the CM might indeed well countenance that all intentional experiences, including those that are about things that do not exist, are about something in the sense of having a representational content. Nonetheless, it must reject the claim that all intentional experiences are about something in the sense of there always being something that one is intentionally directed upon. For instance, if John judges that Pegasus is a horse, then his judgment has a representational content but is not literally about anything since Pegasus does not exist. In such a case, then there is then nothing that is present to John's mind.

Disambiguation often constitutes a helpful philosophical tool. For instance, imagine that someone is trying to analyze the sentence "You have no reason to believe this." Telling that individual that the word "reason" could actually be ambiguous between, say, "epistemic reason" and "practical reason" might turn out to be helpful. For instance, that person might come to realize that this sentence could receive different truth-values depending on how the word "reason" is interpreted. For instance, I might have no epistemic reason to believe that the door is closed but still possess a practical reason for believing that it is; on that ground, I may then try to convince myself that the door is closed (see, *e.g.*, Ginet 2001).

However, in the case of the word "reason," the distinction between, say, epistemic and practical reasons does not follow from endorsing a particular account of reasons. It is simply is a fact *about reasons* that they come in different kinds and any theory of reasons should account for that fact. Hence, whatever grounds we might have to disambiguate between different senses of the word "reason," these grounds are independent of any specific theory of reasons.

The case of "about" turning out to be non-univocal according to the CM is very different. Indeed, in light of Crane's above quotation, that there are two senses of aboutness is not a fact about intentional experiences themselves but instead a fact about their analysis by the CM. Indeed, as Richard Cartwright remarks, it is far from clear that there is any reason to believe that "about" as it is routinely used is ambiguous:

> …it is at least disturbing to be told that, when we finally tell our children that Santa Claus does not exist, we say nothing about Santa Claus. Presumably they expect to hear something about him—the truth about him, one way or the other; and it is scarcely believable that the hard facts of semantics force us to disappoint them. Nor is it much consolation (to us or to them) to be told that we say nothing about him in the same sense as that in which we say something about Caesar when we say he crossed the Rubicon; for it is not clear that "about" has an appropriately different sense. (Cartwright 1960: 633)

If Cartwright is right—and we cannot see any reason to doubt that he is —then the CM's distinction between two senses of "about" is simply *ad hoc* and should actually count against the CM providing a correct description of intentional experiences. Moreover, insofar as one's aim is to provide a correct description of intentional experiences, one should reject the CM's account as an ill-conceived analysis of such experiences.

6. Conclusion

In this chapter, we discussed the main relational alternative to the IOM, the CM. According to the CM, intentional experiences are relations to representational content. This relation is not the one of "being about" but, rather, the one of having a certain content. What an intentional

experience is about is then cashed out in terms of the semantic properties of its representational content.

We then examined how the CM is supposed to account for the six main features of intentionality, focusing particularly on non-existence. The CM is able to dissolve the problem of non-existence faced by the IOM, but it is confronted by a problem of non-existence of its own— namely how to determine the truth-conditions of representational contents that are about things that do not exist. As we saw, the CM must choose between some options here. The task ahead is difficult, but it is commonly believed to be realizable.

Finally, we presented two objections to the CM, based on (i) the existence of non-propositional attitudes and (ii) aboutness. According to the first, the CM schema cannot be generally true as some experiences are about things even though they do not have any representational content. According to the second, the CM schema is not strong enough to capture the intuitive notion of aboutness that characterizes intentional experiences. In the next chapter, we turn to a further model of intentionality, namely adverbialism.

Chap. 5: Adverbialism

> It is possible to avoid commitment to merely intentional objects, then, but the price, if such it is, is the adoption of adverbialism about intentionality. It strikes me that recent discussions of intentionality have treated this web of issues rather glibly.
>
> (Kriegel 2008: 86)

1. Introduction

In the two preceding chapters, we discussed a first family of theories of intentionality, namely the relational theories of intentionality. We discussed two different models: the intentional object model (IOM) and the content model (CM). As we saw, the two models present two very different understandings of the idea that the underlying nature of intentionality is that of a relation. However, both raise some concerns. This provides us with justification to move on to an exploration of another family of theories of intentionality, namely non-relational ones.

In this chapter and the next, we shall discuss two such non-relational theories of intentionality: adverbialism and polyadism. As we shall see, both theories are united in their rejection of the claim that the underlying metaphysical nature of intentionality is that of a relation, but they diverge in their positive claim. Polyadism rejects that the choice between intentionality being either a relation or a monadic property is exhaustive. Adverbialism, on the other hand, claims that intentionality is a monadic property.

In this chapter, we focus on adverbialism, as it has recently been defended by Uriah Kriegel (Kriegel 2007, 2008, 2011). According to adverbialism, intentional properties are adverbial modifications of consciousness. Just as to run quickly is an adverbial modification of a run—that is, a particular way of performing the action of running—to be the bearer of an intentional experience that is about Obama is analyzed by adverbialism as an adverbial modification of consciousness, *i.e.*, as a particular way of being conscious—thinking Obama-wise.

The chapter is organized as follows. We shall begin by clarifying what is meant by a non-relational theory of intentionality. This task accomplished, we shall introduce the basic elements of adverbialism before discussing how it is supposed to account for the six features of intentionality. We shall then consider a first round of standard objections to adverbialism and consider the way in which Kriegel recently proposed to defend adverbialism against these objections. However, as we shall see, a second round of objections will cast doubt on Kriegel's rejoinders.

2. Non-Relationality

In light of the different reasons put forward in favor of the claim proposed in Chapter 2, the claim that intentionality is not a relation might seem implausible. Hence, before entering into the discussion of specific instances of non-relational theories of intentionality, we shall first make two clarifications regarding the claim that intentionality is not a relation.

The first point to notice is that the claim that intentionality is not a relation is not necessarily intended as a denial of the claim that intentionality might *appear* to be a relation. As such, then, non-relational theories need not disagree with the data, in particular grammatical and phenomenological, that speak in favor of treating intentionality as a relation. Non-intentional theories, however, build on the difficulties

faced by relational theories of intentionality to argue in favor of a distinction between appearance and reality: Intentionality may appear to be a relation, but this does not mean that we must treat it as one.

Brentano expresses himself most clearly on this issue in one of his late essays, "Mental Reference as Distinguished from Relation in the Strict Sense," in which he writes:

> The terminus of the *so-called relation* does not in reality need to exist at all. For this reason, one could doubt whether we really are dealing with something relational here, and not rather with something in certain respects *relation-like* (*Relativen Ähnliches*), something which might therefore be called quasi-relational/relational-ish (*Relativliches*). (Brentano 2015: 272)

There is an important issue of translation concerning this passage that was recently indicated by Kriegel (Kriegel 2016). In the German original, Brentano uses the expression "*Relativen Ähnliches*," which can be translated most literally to "similar to a relation." The standard English translation of this passage, however, translates it to "quasi-relation," which might be taken to imply that Brentano believes that there are two kinds of relations: genuine relations on the one hand, for which all the relata must exist; and quasi-relations on the other, for which it is not required that the relata exist. As Kriegel points out, however, the translation of "*Relativen Ähnliches*" by "quasi-relation" is unwarranted and leads to confusion. A more felicitous translation would be "relation-ish," which implies that intentionality does indeed possess the air of a relation without necessarily implying that it is, as a matter of fact, a relation.

In light of that remark, the spirit and task of non-relational theories of intentionality become clearer. They do not necessarily aim to deny that there are data about the phenomenon of intentionality that

speak in favor of treating it as a relation. What they deny is that these data entail that it is a relation. What we should strive for, ideally, is an analysis of the underlying metaphysical nature of intentionality that would allow us to account for these two facts: (i) intentionality *seems* to be a relation; and (ii) intentionality *is not* a relation. One must then try to save the appearances without turning them into facts.

The second point to notice is that the initial plausibility of non-relational theories of intentionality also rests on distinguishing between two different readings of the claim that intentionality is not a relation: a strong reading and a weak one. A first, strong way to read the claim that intentionality is not a relation would be to claim that intentionality is *never* a relation. This, however, would certainly threaten the claim that intentionality can put us in contact with the world, and would make non-relational theories of intentionality fall pray to the objection of idealism alluded to by Grossmann that we reported in Chapter 2.

A second, weaker reading of the thesis that intentionality is not a relation amounts to taking it as meaning that intentionality is not *constitutively* a relation. This second claim, contrary to the first one, is perfectly compatible with the claim that some instances of intentionality genuinely *are* relations. It denies, however, that being relational is constitutive for something to be an instance of intentionality. The claim that a certain instance of intentionality is a relation can then be challenged without this impinging in any sense either on its intentional nature or on the relational character of *some* instances of intentionality.

That said, although denying that a certain instance of intentionality is relational might not impinge on its claim to be a fully fledged instance of intentionality, it might nonetheless impinge on some of its other features, such as its content being true or veridical, or it being an instance of knowledge, or, more metaphorically, it being an instance of a bridge between the mind and reality. As Kriegel puts it,

The deeper point… is that there is a confusion lurking in the way the "cognitive contact" requirement is often wielded. The requirement need not be satisfied by intentionality as such, only by *successful* intentionality. A *true* thought connects us to the world; but it does so in virtue of being true, not in virtue of being a thought. An accurate auditory experience as of trumpets connects us to trumpets, not in virtue of being an experience as of trumpets, however, but in virtue of being accurate. Admittedly, there must be a bridge between intentional states and reality. But intentionality is the bridged, not the bridging. The bridging is provided by such notions as truth, accuracy, veridicality, and the like. What is true of non-veridical experiences is not that they connect us to the word, but that they *would* connect us to the world if they *were* veridical. (Kriegel 2011: 166)

This is a very substantial point that cuts right into the heart of the metaphysics of intentionality: What do we mean by words like "intentional experience," "thought," and "perceptual experience," and how do we conceive of the conditions of individuation of such mental states? The issue is of course very complex, but we shall briefly examine it.

First of all, one must clearly distinguish between genetic and constitutive issues: between the genetic question, "how did we *end up* being such intentional creatures?"; and the subsequent constitutive question, "what *makes* us such intentional creatures?" It is plausible enough that we might have evolved to be intentional creatures by standing in relation to our environment. This, however, is only a claim about the causal origin of our intentional nature, and is not *de facto* equivalent to a constitutive account of intentionality. Non-relational theories might then claim that although the appeal to relations represents a plausible genetic story, it does not represent a plausible constitutive one.

Brains-in-a-vat scenarios are helpful tools for such discussions. We might, for instance, imagine a counterfactual world where, by pure chance and without any hint of design, the following situation materializes: A brain is maintained in life in a vat of nutrients and is connected to a computer that feeds it experiences. Such a brain-in-a-vat scenario certainly does seem to be metaphysically possible, and it also seems perfectly possible for the brain to undergo experiences that are identical to ours such as, say, the cognitive experience characteristic of thinking about Pegasus. Now, clearly enough, the genetic story as to how this brain-in-a-vat ended up being an intentional creature would be much different from our own. That said, that the genetic story would be different does not prevent the constitutive story as to what makes it the case that this brain-in-a-vat is the bearer of intentional experiences to be in principle completely sharable between it and us. And, moreover, since this brain-in-a-vat does not entertain any relation with its environment— or at least not any relation that would plausibly explain why it undergoes just these experiences—this constitutive story that shall account for what we share with it cannot be the genetic one alluded to above.

Externalists about mental contents have of course denied the plausibility of such scenarios. According to them, the individuation conditions of intentional contents are such that they necessarily depend on relations to the environment. Accordingly, it cannot be the case that we share the same intentional life as a brain-in-a-vat that never had any relation with our environment. As Putnam puts it:

> [The brains in vat] images, words, etc., are qualitatively identical with images, words, etc., which do represent trees in *our world*; but we have already seen… that qualitative similarity to something which represents an object… does not make a thing a representation all by itself. In short, the brains in a vat are not thinking about real trees when they think

"there is a tree in front of me" because there is nothing by virtue of which their thought "tree" represents actual trees. (Putnam 1981: 12-13)

The gist of Kriegel's quote above, however, is that a claim such as Putnam's is far too strong. First of all, Putnam is under-describing the case at hand. He treats contents of consciousness on par with images. The imagistic theory of intentionality, however, is absurd. It claims that we can think about things because we can token mental images that stand in a similarity relation to these things. It is clear, however, that similarity is not enough to explain intentionality, as similarity is a symmetric relation. Hence, if my mental image of Obama is about Obama in virtue of being similar to him, then, by the symmetry of similarity, Obama is about my mental image. But this is absurd! Putnam, hence, is not exactly displaying much charity in saddling his opponent with such a view.

Second, properly described, the case at hand does not concern images but appearances. There is no reason, however, to think that appearances are imagistic—and, moreover, no reason to think that appearances have intentionality in virtue of a relation of similarity. Moreover, we can perfectly recognize that there is a fundamental difference between the brain-in-a-vat and ourselves without thereby having to endorse the claim that there is nothing in common intentionally speaking between it and us. We might perfectly share the same experiences without claiming that sharing these experiences puts us in exactly the same situation. To use Kriegel's expression, our experiences, for instance, can routinely be bridged—granted that they are suitably caused by, or related to, reality—while the brain-in-a-vat experiences cannot.

In its strongest form, however, externalism claims that we cannot share anything experientially speaking with the brain-in-a-vat because the content of an experience does not only depend on but is *constituted* by

what it is about, *i.e.*, by the world. Once the world itself is dropped, then the experience and its content must also be dropped. Such views, however, place a heavy burden on the shoulders of its proponents. First, they must deny that the brain-in-a-vat scenario is possible, which seems unlikely. Second, they must explain away either the similarities between the brain-in-a-vat's experiences and our own, or, under the assumption that no such cases are possible, the difference between our actual cases of successful intentionality and our failed intentional states, such as mistaken beliefs and hallucinations.

The general answer given by such radical theories standardly take the following form: There is no such thing as a similarity between successful and failed intentional states to be explained because, they claim, either one is the bearer of a successful intentional state and one is the bearer of a *genuine* intentional state; or one is the bearer of an altogether, non-intentional, different kind of mental state. Such theories are then commonly dubbed "disjunctivist" insofar as they claim that there is no similarity to be explained, *i.e.*, no common-kind, between successful and failed intentional states. One finds oneself either in a good case and is the bearer of an intentional state, or in a bad case, and is not the bearer of an intentional state.

Note, first of all, that a direct consequence of such views is that they estrange us from our mental life. That is, they move the authority regarding what kinds of mental states we are in from the first-person perspective to the third-person perspective. In these views, there is no need to deny that, from the first-person perspective, one cannot know on which side of the disjunction one finds oneself. It is our embedment to a global situation that decides for us on which side of disjunction we are. And, since we have no other way to access this more global situation than our own, first-personal, one, we will not in fact ever know on which side of the disjunction we are. At best, we can now, transcendentally, that

we sometimes find ourselves on the good side of the disjunction. That is a rather big bullet to bite.

Furthermore, such theories cannot simply rest content with their disjunctive claim. They must in addition provide a theory of both disjuncts—the successful and the failed one. But once the richness of conscious intentional life is recognized, it seems impossible for disjunctivism to claim something other than the following: instances of the first disjunct are phenomenally identical to instances of the second disjunct but they fail to put us in contact with reality. That conclusion, however, is identical to the one put forward by non-relational theories of intentionality and, at this point, it seems that the difference between non-relational theories of intentionality and disjunctivism verges on the terminological.[22]

Of course, disjunctivism would require a much longer and elaborate discussion. For the time being, however, we contend that these two reasons provide us with a *prima facie* justification for the claim that refuting disjunctivism does not amount to a top priority. In summary, then, non-relational theories of intentionality need not deny that intentionality appears to be a relation, nor that some instances of intentionality are relational. What they deny is that relationality is constitutive of intentionality. What they should provide us with, hence, is a theory that is able to conciliate two aspects that seem contradictory: the fact that intentionality appears to be a relation, and the fact that

---

[22]About disjunctivism's difficulty to provide us with an adequate account of hallucinatory visual experiences, Michelle Montague, for instance, writes: "[Mike Martin's] account of hallucination is inadequate at the crucial point because it leaves out phenomenal particularity, which can be just as present in a hallucination as it is in veridical perception. And since Martin takes phenomenal particularity seriously in the case of veridical perceptions, we should expect that he would take it just as seriously in the case of hallucinations." (Montague 2011: 135)

intentionality is not a relation. We shall now see how they aim to achieve this task.

3. The General Idea of Adverbialism

We shall now turn to our first non-relational contender, so-called adverbialism. The fundamental idea of adverbialism as a theory of intentionality is that intentional properties are adverbial modifications of experiences. That is, for one to be an intentional subject is to instantiate the first-order property of being conscious, which then instantiates a second-order property, an adverbial property, as of being conscious in particular way—an *intentional* way.

However, before turning to the idea of intentionality being an adverbial phenomenon, we shall introduce the two important theoretical building blocks of adverbialism. First of all, we shall introduce the idea of *adverbial truth-makers* for adverbial sentences. This shall turn out essential as, according to adverbialism, standard intentional sentences such as: "John thinks about x" or "John thinks that p" are made true by such adverbialist truth-makers. And, secondly, we shall introduce the idea of an *adverbialist paraphrase* of a non-adverbial sentence as, according to adverbialism the first important step before providing an adverbialist truth-maker for intentional sentences actually is to offer a systematic strategy to paraphrase intentional sentences in an adverbial form. That is, that "John thinks about x" can be made true by an adverbialist truth-maker makes sense only on the background that "John thinks about x" can be paraphrased into some adverbial sentence.

3.1 Adverbial Truth-Makers

In this section, we shall introduce the idea of adverbial truth-makers for adverbial sentences, that is, the idea that the proposition expressed by an

adverbial sentence such as, "*a ø*s *F*-wise", possesses the following logical form: a particular, *a*, instantiates a first-order property, "to *ø*", which in turn instantiates a second-order property, the adverbial modification "*F*-wise". Accordingly, the truth-maker of such a sentence would amount to the instantiation by *a* of the property of *ø*-ing, duly adverbially modified *F*-wise.

Let us start with the following example:

(1) John runs.

This sentence can be taken to express a certain proposition: namely, that John runs. This proposition is true, let us assume, just in case John instantiates the property of running. Let us, moreover, call the instantiation of this property by John the *fact* that he is running. (1), we might then say, is true just in case the fact that John instantiates the property of running obtains. In that intuitive sense, the fact that John instantiates the property of running can be said to be the truth-maker of the proposition expressed by (1).

Let us now compare this first example with the following one:

(2) John runs quickly.

Informally, the general idea behind (2) can be put as following. Running is an activity that can be performed in many different ways. One can run slowly, quickly, elegantly, effortlessly, and so on. The question, however, is how these ways of doing something impinge on the nature of less determinate propositions, such as the one expressed by (1), and on the nature of less determinate facts, such as the truth-maker of the proposition expressed by (1). In other words, what is the logical form of the propositions expressed by sentences that describe activities done in a

certain way and what is the nature of their truth-makers? Let us try some options.

A first way one may conceive of the proposition expressed by (2) is as a conjunctive proposition. As such, then, the proposition expressed by (2) would simply be the proposition expressed by (1) conjuncted with another one. This conjunctive proposition can be represented as follows:

(3) Run (John) & Quick (John)

Accordingly, (3)'s truth-maker should be conceived as the fact that John instantiates simultaneously both the properties of running and of being quick. However, one might reasonably doubt that (3) really amounts to the proposition expressed by (3). Indeed, (3) could be true even if (2) were to be false, *e.g.*, if John's quickness was completely unrelated to his running, that is, if John was generally quick at what he does though not running quickly. As such, then, (3) does not seem able to capture the intuition that the (2) describes an instance of a certain activity, *i.e.*, a running, done in a certain way, *i.e.*, quickly.

An alternative would be to replace the idea that (2) expresses a conjunctive proposition by the claim that it expresses a simple proposition constituted out of a complex monadic property, *i.e.*, running-quickly. As such, then, the proposition expressed by (2) would amount to the following:

(4) Run-quickly (John)

The truth-maker of that sentence, moreover, would simply be John's instantiation of this complex property.

Once again, however, it does not seem that this corresponds to the correct interpretation of (2). Indeed, in the analysis of (2) offered by (4), "quickly" is a morphological component of the predicate "running-

quickly" but not a syntactical component. As such, then, the two properties "running" and "running quickly" represent, despite morphological similarities, two different properties and, hence, two different activities. For John to run quickly, however, does not seem for him to instantiate a monadic property that is different than the one he instantiates when he is running, though not quickly. Indeed, running and running quickly do not constitute different kinds of activities on par with, say, running and sitting still. Hence, when John accelerates and starts to run quickly, for instance, the kind of activity he is practicing, running, remains steady. It is only the way he performs this activity that changes. Accordingly, (4) seems to fail to capture the proposition expressed by (2) and the fact that John is instantiating the complex monadic property of "running quickly" does not seem to be the right truth-maker for (2).

The idea of adverbial modification enters the stage at this point. So far, we have attempted two different ways to capture the difference between (1) and (2) merely at a first-order level, that is, at the level of properties of individuals such as John. The first option analyzed (2) as John's instantiation of two distinct first-order properties, "running" and "being quick", while the second one analyzed it as John's instantiation of a single complex first-order property, namely "running-quickly". Each of these strategies, however, turned out to be unsuccessful. In other words, it seems that we cannot make sense of the idea that certain activities, like running, can be performed in certain ways by appealing only to individuals and first-order properties. The starting point of adverbialism is then that in order to account for the idea that activities can be performed in certain ways we need to move one level up and appeal to second-order properties, *i.e.*, properties of properties.

Adverbialism's claim, hence, is that the difference between (1) and (2) is not to be accounted for by a difference in the properties instantiated by John himself but, rather, by the presence of a second-

order property that characterizes John's running in some way. This second-order property is an adverbial modification of a first-order property, the property of running. As such, then, the proposition expressed by (2) represents the instantiation by John of the first-order property of running, adverbially modified quick-wise. Accordingly, (2)'s truth-maker would amount to the complex structure of a particular, John, instantiating a property, "running", itself instantiating a second-order property, the adverbial modification "quick-wise."

These ideas of adverbial modifications, adverbial propositions and adverbial truth-makers introduced, let us now turn to the complementary issue of providing adverbial paraphrases of non-adverbial sentences.

3.2 Adverbial Paraphrases

The gist of adverbialism as a theory of intentionality is that the intentional character of intentional experiences can be accounted for by treating them as having an adverbial nature. As we saw in the above example, to run quickly is for a certain activity—*e.g.*, running—to instantiate a certain adverbial modification. In the same way, to be the bearer of an intentional experience, according to adverbialism, is for an experience to be adverbially modified in some way, *e.g.*, to think about some *x* must be treated as thinking *x*-wise.

There is, however, an important difference between examples such as (3) and intentional experiences. Namely, (3) explicitly is an adverbial sentence, and the question we tackled was how we should interpret such sentences. On the other hand, intentional sentences, *i.e.*, sentences describing intentional experiences, are not adverbial sentences or, at best, not explicitly. An important task should then be to tackle the question of how one could provide an adverbial paraphrase of a non-adverbial sentence.

As we shall see, the use of adverbial paraphrases to avoid metaphysical conundrum is a familiar philosophical move that is not unique to the theory of intentionality, and other kinds of problematic sentences can be submitted to an adverbialist treatment to avoid some problematic commitments. We shall then take our cue from there before moving to intentional sentences.

In his discussion of intentionalism about perception, William Fish (Fish 2010: 34), for instance, discusses the following sentence:

(6) John wore a broad smile.

The surface grammar of this sentence is relational. It seems to express a relation, "wearing (x,y)", that obtains between two things: John and a broad smile. In such a reading, however, the indefinite noun-phrase "a broad smile" is supposedly referential, as if John ever wore a broad smile, he definitely wore a definite one. That said, the idea that there could be determinate smiles to which we could stand in relation turns out to be rather unpleasant. Indeed, what kinds of particulars are smiles and how should we individuate them? But if (6) indeed has a relational structure, then we would be forced to accept that there thus are smiles to which we can stand in relation and to which indefinites such as "a smile" would refer in order to account for the truth of (6) understood as expressing a relational proposition.

In the face of this difficulty, a popular alternative is to claim that (6) is true, but only non-literally so. That is, it is true because a paraphrase of (6) is true. This paraphrase, moreover, is an adverbial one:

(7) John smiled broadly.

(7), contrary to (6), does not feature a relational surface grammar. Rather, it can be parsed as the instantiation by John of a first-order property,

"smile (x)." We can then treat the adverb "broadly" in the same way in which we treated "quickly" in the case of (3)—to smile broadly is simply to smile in a particular way, to be accounted for in terms of the presence of a second-order adverbial modification of John's smiling. By means of such a treatment, all concerns regarding whether we could stand in relation to smiles evaporate and all concerns regarding a non-adverbial fact that (6) would describe disappear. What is actually meant by (6), the claim goes, is that John indeed smiled, though he did so in a particular way, namely broadly.

An important fact that we shall bear in mind for the upcoming discussion of adverbialism is how natural the move from (6) to (7) seems to be. Indeed, tested against the intuitions of native speakers, it is very likely that (6) and (7) would be taken to mean the same thing. In fact, one could even imagine that (7) be taken to reveal the meaning of (6) to some benighted speakers. Indeed, upon being told (6), one could be unaware of the use of the verb "to wear" in such a non-literal context and imagine that (6) describes John wearing a huge manufactured smile pinned on his sleeve. That is, after all, the literal meaning of the verb "to wear." (7), then, can be said to provide a possibly welcome disambiguation of (6) and, hence, to be a way to reveal its proper meaning in most usual contexts.

We shall remain very general as to the strategy that one should use in order to turn a non-adverbial sentence into an adverbial one, preferring examples to systematic guidelines. However, one key to performing such transformations, as illustrated by the passage from (6) to (7), seems to be the replacement of a transitive verb, such as "to wear," by an intransitive verb and the transformation of an adjective, such as "broad," into an adverb. In the next section, we shall see whether these elements of paraphrasing can be applied as such to intentional sentences, or whether some further means must be appealed to.

4. Adverbial Intentionality

Adverbialism about intentionality aims, just as the CM does, to provide a paraphrastic interpretation of the FIS. An adverbial interpretation of the FIS is one that substantially departs from the original schema by turning it into an adverbial construction from which the little word "about" has disappeared. The metaphysical side of this idea is then that intentional sentences paraphrased in adverbialist terms possess truth-makers that feature a metaphysical structure similar to the one discussed above in the cases of (2) and (5). In this section, we shall explain the main elements of such an account of intentionality.

Let us begin with a standard intentional sentence, such as the following:

(8) John thinks about Obama.

This sentence is, of course, not an adverbial one, and the first adverbialist step should be to provide an adverbialist paraphrase of (8). The general idea is that we shall turn (8), which seems to describe a relational fact obtaining between two things, John and Obama, into an activity of John's that he can perform in a certain way.

According to adverbialism, this activity of John's should, most basically, be conscious. There are, however, different kinds of conscious states, such as sensations, and more cognitive activities, such as thinking. Thus, adverbialism proposes that the kind of activity with which John is busy when he is thinking about Obama is not merely being conscious, but rather instantiating a particular kind of conscious property, a thinking property. The first part of our paraphrase, then, is to turn the transitive "to think about something" into the intransitive "to think." Thus, according to adverbialism, when John thinks about Obama he first and foremost thinks.

However, if (8) is true, then John is not merely thinking, but rather thinking about Obama. Here comes the second step of the adverbialist transformation: the creation of an adverbialist determination from the remaining parts of (8). In the previous sub-section, we turned an adjective into an adverb but this time we should turn a noun, Obama, into an adverb, *i.e.*, Obama-ly or, more elegantly, Obama-wise.

Putting these two elements together, the adverbialist treatment of a sentence such as (8) amounts to the following:

(9) John thinks Obama-wise.

At this point, one might have an initial concern. Indeed, one might point to the fact that in the case of the transition from (6) to (7), the use of "smile" is unproblematic because "to smile" not only is an intransitive construction, but is also a construction that does not imply the truth of a transitive one. The case, however, seems to be different with "to think." Indeed, one can merely smile, but one cannot merely think; to think is an intentional activity. Hence, one cannot merely think but must always think about something. As such, then, one might worry that the transition from (8) to its adverbial paraphrase (9) does not run exactly as smoothly as the transition from (6) to (7).

This concern, however, is misplaced. Indeed, it seems to rely, first of all, on the idea that we have a genuine grasp on intentionality that is independent of adverbialism. But this is precisely what adverbialism argues that we do not have. Adverbialism is a theory of intentionality and, if it is true, then it is simply false that there is no reading of the schema "*x* thinks" that does not entail an instance of the transitive one, "*x* thinks about *y*." Rather, according to adverbialism, the nature of intentionality would force us to say that there is no reading of the schema "*x* thinks" that does not entail an instance of the adverbialist

schema "*x* thinks in some way." As such, then, this first concern simply seems to be begging the question with respect to adverbialism.

Moreover, one might also point to the fact that the sentence, "One can merely smile," is ambiguous between two readings. It can be interpreted as meaning that the verb "to smile" does not imply a transitive construction, in which case it is correct, but it can also be interpreted in the sense that one can smile without smiling in some way, in which case it is simply mistaken. Indeed, one cannot smile without smiling in some way, *i.e.*, broadly, shyly, expressively, ironically, and so forth. But from the point of view of adverbialism, "to smile" is then very much like "to think."

Let us grant for the sake of the argument that one can indeed successfully paraphrase intentional sentences such as (8) into adverbial ones such as (9). That said, one might still wonder what exactly happens beneath the mere linguistic transition—that is, how we should understand the transition from the side not of the sentences themselves, but rather of what they are taken to describe, *i.e.*, intentional experiences.

Simply put, adverbialism's claim is that intentional experiences have an adverbial nature. In other words, they are constituted by the instantiation of a certain first-order property—namely being conscious or, more determinately, thinking, conceived as a conscious property— which is then adverbially modified. Imagine, for instance, that you are the bearer of a cognitive experience such as thinking about chocolate. According to adverbialism, this is to be understood as being conscious in a certain way. Being conscious is a first-order property of yours; whereas thinking chocolate-wise is a second-order, adverbial modification of that first-order property. The point, according to Kriegel, can be generalized to all instances of intentionality. For each kind of intentional state, there is an adverbial modification that corresponds to that state. In general terms, if one is the bearer of an intentional state *S*, then one is intentionally directed *F*-wise. As Kriegel puts it,

> According to this adverbial theory, having a conscious thought about
> chocolate, say, does not involve (constitutively) bearing a thinking-about
> relation to the property of being chocolate, but rather engaging in the
> activity of thinking in a certain *way*–chocolate-wise. Just as in "Omar is
> moving quietly" the adverb "quietly" denotes a modification of Omar's
> moving, and not a relation between Omar and quiet, so in "Omar is
> thinking chocolate-wise" the adverb "chocolate-wise" denotes a
> modification of Omar's thinking, not a relation between Omar and
> chocolate. More generally, for any putative property F, for F to figure in
> the exp-intentional content of a state x is not for x to bear a relation of
> intentional directedness toward F; rather, it is for x to exhibit the
> property of intentional directedness F-wise. (Kriegel 2011: 153)

The difference between (8) and (9) is then completely parallel to the one
between (6) and (7). (8) has a relational structure, while (9) does not. It
consists of the predication of the property "think (x)" by John. The
adverb "Obama-wise" can then be taken as modifying the verb "to
think" in the lines described above. As such, then, a sentence like (9)
does not require as a truth-maker a relational fact of the kind, "Thinks
about (John, Obama)." Instead, it can be true if the non-relational fact
that John's conscious state of thinking instantiates the second-order
property of thinking Obama-wise. As such, then, adverbialism treats the
metaphysics of intentional experiences in a completely non-relational
manner, and without presupposing at a fundamental level any of the two
key notions that we discussed in the context of the IOM and the CM—
the one of intentional object and the one of representational content.

At this point, however, one might well wonder how come exactly
we know that intentional experiences have such an adverbial nature.
According to Kriegel, the plausibility of adverbialism as a theory of
intentionality relies heavily on realizing that the inference to adverbialism

is of an *adbuctive* nature. That is, from the point of view of adverbialism, we have access to our intentional experiences from our first-personal perspective and these experiences have, as a matter of fact, an adverbial nature. But this does not mean that we have direct access to their adverbial nature. On the contrary, that the nature of these experiences is in fact adverbial is something that we have to infer later on theoretical grounds. As Kriegel puts it,

> When I reflect on why adverbialism seems *to me* perfectly intelligible, and how *I* get a sense of grasping the property of being directed dragon-wise, it seems to be that it is simply introspective encounter with the property of being directed dragon-wise that affords me this grasp. By this I do not mean that experience instructs me of the adverbial nature of exp-intentionality; merely that it instructs me of what in fact are adverbial exp-intentional properties. Thus, I just now stopped writing and visualized a smallish (about a foot long) green dragon hovering motionless about a yard away from me directly in front of my eyes. When I attended introspectively to this visualizing experience, I encountered a certain property of the experience, a property we have theoretical reasons for construing as visualizing smallish-green-hovering-dragon-wise. (Kriegel 2011: 155)

The argument for adverbialism therefore comes in two steps. First, we have introspective access to our intentional experiences and to their intentional nature, and second, we have theoretical reasons to construct their nature adverbially. There is no direct access to their adverbial nature. We say more below, however, about the nature of the abductive inference that concludes that adverbialism should be the way to go for a theory of intentionality.

Having presented the great lines of the adverbialist proposal, we shall, in the next section, see how adverbialism can be said to be able to

account for the six main features of intentionality that we identified in Chapter 1.

5. Adverbialism and the Seven Features of Intentionality
5.1 Aboutness

According to adverbialism, aboutness is a non-relational property of an experience, namely the property that Kriegel calls "intentional directedness *F*-wise." That is, for an experience to be intentional is for it to instantiate a particular kind of adverbial modification, an intentional one, that turns it into an intentional experience. As such, then, we can say that adverbialism proposes a non-literal interpretation of the FIS that goes the following way:

> **Adverbialist schema**: "*x* is about *y*" amounts to "*x* is intentionally directed *y*-wise."

There are a few points to be made about this schema. A first one is that it seems rather intuitive that at least some intentional experiences are, as Kriegel puts it, "object involving" (Kriegel 2011: 163) or, as they as standardly called, "singular." The intuitive idea here is that intentional experiences that are about particulars, such as (9), are such that they involve the direct presence of a particular to the mind and not merely the presence of properties of this particular.

Different models are likely to account for this claim in different ways. The IOM, first, will account for this claim by saying that in such a case, the relatum of a such an "object-involving" experience is the particular it is about. The CM, on the other hand, traditionally accounts for this distinction by relying on the distinction we discussed in Chapter 4 between so-called general and singular truth-conditions. That is, experiences that are object-involving are said to have as representational

content a proposition that has singular truth-conditions, that is, a proposition whose very existence and identity depends on the particular it is about. While experiences that are not object involving, on the other hand, have as content mere specifications of what they are about. Moreover, though there recently has been some debate about this issue (*e.g.*, Jeshion 2010a), the CM generally claims that for one to be able to token such singular representational content, one must be acquainted in some way or other with the object upon which a singular proposition that serves as representational content existentially depends.

That said, independently of the way this intuitive distinction is implemented, the idea is that there is a substantial difference between these experiences that are object-involving and those that are not. This can either be accounted for in terms of a difference in the kind of relata one is intentionally directed upon, as in the IOM, or, as in the CM, in the kind of representational content a certain intentional experience is having.

Adverbialism, however, does not possess the means to account for this distinction. Indeed, either the IOM or the CM requires that such singular experiences be individuated in terms of a relation to particulars. According to adverbialism, however, there is no instance of experiences that, *qua* intentional, are thus individuated. At best, some intentional experiences can, as Kriegel puts it, "be bridged" (Kriegel 2011: 166) but this does not mean that the contents of any experience is individuated in an object-involving way.

According to Kriegel, however, the fact that adverbialism is unable to account for the distinction between intentional experiences that are object-involving and intentional experiences that are not cannot serve as an objection against adverbialism. First, he points to a consequence of the distinction that he is unwilling to endorse. Second, he proposes to disentangle two different issues, namely: (i) the idea that there is a difference between intentional experiences that are about particulars as

well as properties and intentional experiences that merely are about properties and (ii) the idea that we must appeal to the distinction between object-involving and non-object-involving experiences to account for this distinction. As he sees the matter, (i) can be accounted for independently of (ii), in terms of a distinction between two features, generality and "particularity" (Kriegel 2011: 163).

The first point is the following. A consequence of the claim that some experiences are object-involving is that experiences that are about numerically distinct, phenomenally indistinguishable particulars will have different representational contents, as these contents will be individuated in terms of the objects they are about, regardless of the fact that they cannot be distinguished phenomenologically. According to Kriegel, however, this result is unintuitive. As he puts it:

> Let Twin-World be a possible world indistinguishable from the actual world in every detail but one: the laptop I am looking at right now in that world is a numerically distinct individual from the one I am looking at in the actual world. Although qualitatively indistinguishable from any actual laptop, it is *another* laptop (perhaps because it has a different *haecceity*, perhaps for some other reason—depending on one's view of what makes a concrete particular the concrete particular it is). According to the objector [*i.e.*, the proponent of a distinction between object-involving and non-object-involving experiences], the exp-content of my experience in the actual world and in Twin-World is different. This verdict, however, strikes me as obviously counter-intuitive. It is much more intuitive that the exp-intentional content are the same—that the two experiences are exp-intentionally type-identical. (Kriegel 2011: 163-164)

In other words, according to Kriegel, the distinction between object-involving and non-object-involving experiences cannot serve as an objection against adverbialism because it is wedded to a so-called

"externalist" individuation of the content of intentional experiences[23]. That is, it is wedded to an account of the individuations of intentional experiences according to which the crucial facts for the individuation of representational contents are not how the world appears to a subject but, rather, facts that lie beyond such appearances. Accordingly, Kriegel's thought about his laptop in the actual world and about his laptop in Twin-World have different contents because, despite the fact that the world appears to Kriegel as being the same in both cases, the two thoughts in fact turn out to be about different individuals. According to Kriegel, however, externalism is implausible as it is, as he puts it, "obviously counter-intuitive" that the two experiences have different contents. Accordingly, all contents of intentional experiences have, as they are called, internalist individuation conditions, *i.e.*, individuation conditions that do not tie them to the objects they are about.

The issue over externalism and internalism is a complex one, central to contemporary debates in the philosophy of mind. Kriegel, however, is not alone in defending internalist conditions of individuation for the content of intentional experiences (see, *inter alia*, Crane 2001; Horgan and Tienson 2002; Farkas 2008). As such, then, we should discuss further the issue it but merely flag Kriegel endorsement of internalism.

Let us now turn to the second issue, namely the general intuition that lies behind the distinction between object-involving and non-object-involving intentional experiences. What, first, is this general intuition that lies behind the distinction? Contrast the following two things: First, being appeared by a manifold of properties and, second, being presented with a particular as through a manifold of properties. The first experience would be a purely general experience of a certain distribution

---

[23] For a useful collection of papers on externalism, see Goldberg and Pessin 1996.

of properties, the second, however, would be an experience as of being appeared with a particular, as it appears to us through a certain manifold of properties. The contrast between these two things, we shall contend, is the general intuition that lies behind the distinction between object-involving and non-object-involving intentional experiences.

According to Kriegel, however, we do not need to appeal to the notion of singularity to account for the contrast between these two kinds of experiences. Rather, he claims, we can account for it in terms of the phenomenological notion of "particularity", that is, the particular experiential character of experiences that present us with particulars. Drawing on work by Michelle Montague (Montague 2011), he cashes out this idea as the fact that some experiences have an "object positing phenomenology." Montague's general idea is that intentional experiences as of being appeared with particulars, rather than mere manifolds of properties, are such that they possess a thought-component that structures the experience into an experience as of being presented with an object. It is this added layer of thinking that turn general intentional experiences into ones that instead are *particular*. Kriegel endorses Montague's idea but proposes, in addition, to turn it into an adverbial feature of experiences. That is, for one to undergo an intentional experience as of being appeared with a particular is for one to be the bearer of an intentional experience that is intentionally directed someway "concrete-particular-wise." As he puts it,

> Although I deny that exp-intentional content is characterized by *singularity*, I recognize that it is characterized by *particularity*, in the sense that its bearers often have the experiential character of presenting concrete particulars. This can be accomodated by claiming that such experiential characters exhibit…an *object-positing phenomenology*, which is the phenomenology of particularity, and then offer an adverbial gloss on this phenomenology. The adverbial gloss construes the phenomenology of

particularity as the experiential character of being intentionally directed concrete-particular-wise. (Kriegel 2011: 164)

As such, then, Kripke's adverbialist interpretation of the FIS is in fact more complex than the above adverbialist schema. Namely, according to him, we need to break down further the property of "being intentionally directed *x*-wise." Take, for instance, the case of John thinking about Obama. According to Kriegel, this will not simply mean that John is intentionally directed Obama-wise. Rather, it will mean that John's experience possesses a complex adverbial content composed of (i) properties of an experience as of being presented with Obama and (ii) particularity. This clarification made, we shall actually stick to the above formulation of the adverbialist schema, duly flagging the fact that Kriegel breaks down further the "*y*-wise" part.

A second point is that the adverbialist schema is, just like the CM schema, a reductive schema in the sense that it tries to give an account of the meaning of the FIS that proceeds without the word "about." Indeed, "about" has completely disappeared from the right-hand side of the adverbialist schema. That said, the adverbialist schema is very different from the CM schema is the sense that it does not try to derive the notion of aboutness from a more fundamental feature, such as semantic normativity. Indeed, aboutness is taken by adverbialism as a fundamental feature of intentional experiences, one that needs to be understood in terms of the adverbial notion of intentional directedness.

Kriegel is quite clear about this point. Indeed, according to him, the notion of directedness should be taken as sufficient to capture the notion of aboutness. This, he thinks, can be made clearer by means of the following two claims that, he says, "could indirectly shed light on *F*-wise directedness." (Kriegel 2011: 154)

(A) Whenever an intentional state of being directed *F*-wise is veridical, then it stands in a relation to some *F*;

(B) Whenever an intentional state of being directed *F*-wise is non-veridical, it is still true of it that if it *were* veridical, it *would* be related to some *F*.

That is, according to Kriegel, there is a quite clear semantic connection, either actual or counterfactual, between an adverbial modification, such as being directed *F*-wise, and some *F*. However, this connection is not supposed to reveal anything about the nature of aboutness itself, as in the CM. Indeed, adverbialism does not aim to discharge aboutness in terms of the truth of the claims (A) and (B). On the contrary, according to adverbialism, we have a grasp on the aboutness of intentional experiences that is *independent* of the truth of (A) and (B), in terms of the notion of "being intentionally directed *F*-wise." But the truth of (A) and (B) is supposed to show us that this notion is *sufficient* to capture aboutness.

A third point is that the disappearance of the word "about" from the right-hand side of the adverbialist schema, and this despite the mention of the above two claims (A) and (B), raises some question about the nature of the paraphrase of the FIS offered by adverbialism. Indeed, how should we evaluate the examples of adverbialist paraphrases mentioned above, as well as the generalization of the adverbialist paraphrasing strategy materialized in the adverbialist schema? Is, for instance, the above paraphrase, from

(8) John thinks about Obama, to

(9) John thinks Obama-wise,

as perspicuous as the one from

(10) John wore a broad smile, to

(11) John smiles broadly?

Before answering that question, let us remember why we should bother to ask it. In accordance with the policy introduced in Chapter 3, an acceptable paraphrase of the FIS should be a conservative paraphrase, that is, a paraphrase whose aim it is to uncover the meaning of natural language sentences and that is accountable to native speakers' intuitions. This certainly is the case for (6) and (7). But is it the case for (8) and (9)?

One concern that one might put forward is that the meaning of (9) taken in isolation is too thin for the theory of intentionality. That is, all things considered, we understand the meaning of (9) as being an adverbialist interpretation of an intentional sentence only because we have an independent grasp on intentionality through sentences like (8). In other words, it seems that the intentional reading of a sentence like (9) is parasitic on a sentence like (8), in which case (9) cannot be said to really provide us with the meaning of (8) in a sense required by a conservative conception of paraphrases.

This charge might turn out to be too quick insofar as the adverbialist has not yet been given the opportunity to elaborate further on her adverbial interpretation of the FIS. For the time being, however, it marks a query to which adverbialism should be able to answer.

That said, all adverbialist that he is, Kriegel seems to be sensitive to the query. Upon being confronted with this objection in one of his early papers on adverbialism, he writes:

> My response to this objection is twofold. First, if the adverbial paraphrase were revisionary, it would still be quite justified. The alternative, after all, is to plunge into the can of worms of merely intentional objects. If our everyday intentional talk is relational rather

than adverbial, then our mundane conceptual scheme is committed to the idea that we routinely think of and perceive objects that either do not exist, or are abstract, or are mental. Any of these commitments, it seems to me, is so embarrassing as to invite radicalism. It is quite probable that even proponents of merely intentional objects accept the following conditional: if we can avoid commitment to merely intentional objects, we should. It is just that they deny the antecedent. (Kriegel 2008: 91)

Kriegel therefore seems to concede the claim that adverbial paraphrases of intentional sentences might not pass the test for the qualification of conservative paraphrase—*i.e.*, that (9) might not provide us with the meaning of (8) in the same way that (7) can be said to provide us with the meaning of (6). His answer, however, is that even if adverbial paraphrases counted as revisionary, we would still be justified in endorsing adverbialism by a kind of argument by elimination: all of the other options are so bad that, all things considered, a revisionist adverbialism does not turn out to be *that* bad.

Kriegel's remark, however, puts him in a complex dialectic. Indeed, first of all, one might challenge the antecedent of his claim, namely that there is no way to discharge the commitments of the IOM in an acceptable way or, as he puts it, that "any of these commitments… is so embarrassing as to invite radicalism." If one has in mind a version of the IOM coupled with a version of Meinongianism that corresponds to the one discussed in Chapter 3, then Kriegel's antecedent might indeed have the force that he thinks it has. That said, however, as we shall see in the next chapter, there are other ways to discharge the commitments of the IOM than the ones we elaborated in Chapter 3, and it is not so clear that they are as embarrassing as Kriegel himself believes them to be. Second, granted that Kriegel is right about the embarrassment of the IOM, his claim would go through only if adverbialism does not find itself with an amount of embarrassment

equivalent to the IOM. As we shall see later on, however, it is far from clear that adverbialism can avoid such a situation.

That said, in his later writings on adverbialism, Kriegel provides an analysis of the property of being directed *F*-wise that might be interpreted as an attempt to make our understanding of an adverbialist paraphrase like (9) be less dependent on an intentional sentence like (8) for its intentional reading.

This problem, Kriegel claims, "calls for a phenomenological analysis of, or description, of the property of being directed *somehow.*" (Kriegel 2011: 156) According to him, the best account that we possess of that notion is in terms of the notion of "phenomenal foreignness" (Kriegel 2011: 158). Simply put, the idea is the following. When someone like John is thinking about Pegasus, John is thinking Pegasus-wise and experiences this adverbial modification as foreign. In other words, an intentional experience with an adverbial nature is an intrinsic modification of the consciousness of a subject that presents something as foreign. The notion of "aboutness," in other words, is cashed out in terms of the notion of "foreignness." An intentional state is about whatever, in an intentional experience, is presented as phenomenally foreign.

That Kriegel's adverbialism explicitly harbors this feature makes it, from the point of view of the description of intentional experiences, intrinsically superior to the CM. Indeed, in the case of the CM, we saw that we had some difficulties in recollecting the intuitive sense of aboutness that we deemed to be captured by the universal application of Existential Generalization (or its Meinongian Positive Free Logic equivalent). One might claim, however, that the mention of phenomenal foreignness is as close as it gets to capturing this feeling of being presented with something that the application of EG is claimed to capture so well, but without presupposing that EG can be applied along the lines defended by the IOM. Indeed, one might claim, talk of being

mentally directed upon an intentional object can then be paraphrased as talk of being presented in some way "foreignly," where this is supposed to be able to account for this feeling of being mentally directed upon something. This might even be regarded as a way to discharge Brentano's expression of something being "relation-ish," mentioned earlier in this chapter.

That said, one might wonder whether this really sets adverbialism apart from the CM. Indeed, one could perfectly defend an augmented version of the CM according to which all instances of intentionality would, in addition to instantiating content, also feature phenomenal foreignness. Accordingly, the content would always be taken as presenting something foreign.

This claim, however, does not impinge on the one that adverbialism might have other advantages over the CM. One of them might simply be the following. Imagine that one is thinking about an impossible property, F. According to the CM, this would mean that F would have to be featured in the content of one's intentional experience. Hence, one would be committed to impossible properties, such as F, a commitment that one may well prefer to avoid. Adverbialism, on the other hand, would not share such a commitment. Indeed, granted that "being intentionally directed F-wise" is not parasitic on the property F itself, then adverbialism could account for the intentional state being directed at impossible properties without being committed to them.

5.2 Adverbialism and Aspectuality

According to aspectuality, all instances of intentionality are instances of thinking about something in some way. Adverbialism does seem to possess the resources to account for this feature. Indeed, there is no reason to deny that different intentional experiences that instantiate different adverbial modifications can in fact be about the same thing.

5.3 Adverbialism and Semantic Normativity

Let us take an example of an intentional experience that would feature semantic normativity such as, for instance,

(12) John judges that the morning star is a planet.

The IOM and the CM agree on distinguishing two important features of an intentional experience like (12): its mode and its content. Adverbialism can also proceed to such a distinction, but in its own way. That is, to judge will correspond to a particular kind of conscious experience, with a particular judging character. This experience, moreover, will be further determined as instantiating some further adverbial modifications, like judging morning-wise, star-wise, and planet-wise. Hence, (12) will turn out to be true just in case there really is something that corresponds to the adverbial content of that state.

5.4 Adverbialism and Intentional Identity

Adverbialism will presumably have to propose two different notions of intentional identity: one for intentional experiences that are about things that exist, and one for intentional experiences that are about things that do not exist. Let us begin with the latter case. Since, just like in the CM, there is nothing, literally, that one is thinking about when undergoing an experiences such as being intentionally directed upon something that does not exist, then there is not anything that two distinct intentional subjects can be said to think about. That said, these two intentional subjects could nonetheless be said to be intentionally identical in another, qualitative sense. Indeed, they could be said to both instantiate the same adverbial modification.

When, on the other hand, the question concerns intentional experiences that are about things that exist, then the truth of the above (A) should ensure that different subjects instantiating different tokens of the same adverbial modification turn out to be thinking about the same thing.

5.5 Adverbialism and Generality

Generality does not seem to create much of a difficulty for adverbialism. Indeed, if John is thinking about a man, then John is merely thinking man-wise—end of story. If adverbialism has some issues, they certainly are not with generality.

5.6 Adverbialism and Non-Existence

According to Kriegel, non-existence constitutes the main reason that we possess to endorse adverbialism. That is, as he made clear in a passage already quoted above, all of the other options are so bad that adverbialism should appear as a natural option (*cf.*, Kriegel 2008: 91). And, indeed, non-existence does not seem to create much of a problem for adverbialism as there is no reason to think that for one to be intentionally directed F-wise, F must exist. That there is an adverbial modification of some kind does not entail that something of some kind exists that corresponds to this adverbial modification.

6. Objections to Adverbialism, Round 1

Kriegel recently went out of his way to defend an adverbialist theory of intentionality, but he did not discover the theory. Adverbialism was indeed much discussed, and objected to, in the 1960s and 1970s, essentially under the form of a theory of perception. These standard

objections, however, carry over adverbialism taken as a theory of intentionality as whole. Hence, in this section we present two main standard objections to adverbialism, as well as Kriegel's general rejoinder to both of them. In the next section, we then turn to objections directed at Kriegel's updated version of adverbialism.

6.1 Adverbialism and Inferences

Frank Jackson coined a famous objection against the adverbial theory of perception (Jackon 1977). This objection, however, can be expended and treated as a generalized objection against an adverbial theory of intentionality along the lines proposed by Kriegel in 2011. The argument has two steps. The first step ends with the intermediary conclusion that adverbialism cannot account for basic inferences that a theory of intentionality should be able to countenance. The second step ends with the conclusion that a slightly modified account of adverbialism can account for these basic inferential patterns, but only at the price of a damageable loss of compositionality. The final conclusion, then, is that adverbialism should be rejected.

The first part of the argument starts with the following premise:

(1) The following two intentional experiences are not equivalent:

    a) I am thinking about a green dragon and a purple
       butterfly; and
    b) I am thinking about a purple dragon and a green
       butterfly.

This first premise does not need discussion and can simply be assumed. Let us now turn to the second premise:

(2) Adverbialism is not able to account for the non-equivalence of
a) and b);

Unlike (1), this premise needs to be supported by means of an argument, along the following lines. A basic version of adverbialism must construe a) as follows:

a') I think dragon-wise, butterfly-wise, green-wise, and purple-wise.

The problem, however, is that adverbialism is also bound to give the same answer in the case of b). Hence, adverbialism is not able to account for the truth of (1).

From (1) and (2), we can then conclude the following:

(3) Adverbialism is false.

The derivation of this conclusion concludes the first step of Jackson's objection. Adverbialism cannot account for the distinction between a) and b) and, hence, should be rejected. The second step then begins with the observation that the treatment of a) and b) offered above is not forced upon adverbialism. Indeed, one could instead offer a more sophisticated version of adverbialism that would be able to offer non-equivalent adverbialist paraphrases of a) and b).

This constitutes the gist of the next premise:

(4) A more sophisticated form of adverbialism does possess the means to paraphrase a) and b) in   such a way that they turn out to be non-equivalent, namely as follows:

a") I think green-dragon-wise and purple-butterfly-wise;

  and:

b') I think purple-dragon-wise and green-butterfly-wise.

The two paraphrases a') and a") are in fact very different. Indeed, a') was constructed out of one first-order predicate, thinking, and four second-order predicates: dragon-wise, butterfly-wise, green-wise, and purple-wise. On the other hand, a") is composed of one first-order predicate, thinking, and two second-order predicates, green-dragon-wise and purple-butterfly-wise; and so is b'), although it is composed of different second-order predicates, namely purple-dragon-wise and green-butterfly-wise.

One may then think that the adding of (4) paves the way for a successful version of adverbialism, as its endorsement would be sufficient to account for the truth of (1). The problem, however, is that the endorsement of (4) clashes with other requirements for which an acceptable theory of intentionality should be able to account. Take, for instance, the following inference:

(i) I am thinking about a green dragon; and

(ii) Hence, I am thinking about a dragon.

This inference is enthymematic but it is certainly valid once we reconstruct it in a more explicit form, such as in:

(iii) I am thinking about something that is a dragon and that is green; and:

(iv) Hence, I am thinking about something that is a dragon.

Indeed, reconstructed under this form, this inference looks like a perfectly valid instance of conjunction elimination, and we can then formulate the following fifth premise:

> (5) An acceptable theory of intentionality should be able to account for some basic inferences, such as the above one going from (iii) to (iv).

The problem for adverbialism, however, is that moving away from its naive formulation to a more sophisticated one in order to avoid the problem of distinguishing between a) and b), it is now impaling itself on another problem: being unable to account for the inference from (iii) to (iv). Indeed, from

> (v) I am thinking green-dragon-wise,

it is impossible to appeal to a principle like conjunction elimination in order to validly infer the following:

> (vi) Hence, I am thinking dragon-wise.

Indeed, it is essential to point to the fact that the second-order predicates newly formed by the sophisticated version of adverbialism appealed to in premise (4) have morphological but *not* syntactic components. That is, in the predicate "green-dragon-wise," "green" occurs as a morphological and not a syntactical component. As such, then, the predicates "green-dragon-wise" and "purple-dragon-wise," for instance, have as much in common as the two words "leg" and "legitimate," which both share the string of letters "l-e-g" as a morphological component. From this, then, we can infer the following:

(6) Sophisticated adverbialism cannot account for the validity of
    the inference from (i) and (ii);

And, finally, from (6) one can simply conclude that adverbialism is mistaken.

Does this objection constitute a definite rebuttal of adverbialism? Kriegel (2011) offers a rejoinder to this first objection, based on the familiar distinction between determinables and determinates. However, as this answer is also the one that Kriegel offers to the second objection, we postpone its discussion and present first the second standard objection to adverbialism taken from Dermot Moran's work.

6.2 The Type-Similarity Objection

In his paper "Brentano's Thesis," Moran offers a quick dismissal of adverbialism along the following lines:

> Briefly, if intentional objects are to be construed adverbially in this manner, the danger is that all acts would be quite distinct from each other in kind, infinitely multiplying mental acts. (Moran 1996: 9)

This objection is related, but distinct, from the previous one. It is related because it takes as its starting point the above premise (4). It is distinct, however, as it is not directly concerned with the inferential properties of some intentional sentences, but rather with the type-similarity of intentional experiences themselves. The basic idea is the following. Intuitively, thinking about Obama and thinking about the Queen mother have something in common: They are both instances of the same kind of intentional experiences by virtue of sharing what we may call a same mode or attitude, namely thinking. As such, then, they can be typed

along in a way that cannot be shared by an instance of, say, perceiving Obama.

The problem for the sophisticated version of adverbialism, however, is that it is not able to account for these intuitive type-similarities. Indeed, since, once again, "thinking-Obama-wise" and "thinking-Queen-Mother-wise" have no syntactic component and constitute two completely distinct adverbial modifications, there is no basis to support the type-identification of these two states and no reason to claim that "thinking about Obama" and "thinking about the Queen Mother" share something that they do not share with, say, "perceiving Obama." From this, one may then conclude, once again, that adverbialism is mistaken.

6.3 Kriegel's Rejoinder to the Two First Objections

In *The Sources of Intentionality* (Kriegel 2011), Kriegel proposes a rejoinder to Jackon's objection that can also be applied to the second objection. Kriegel's aim is to find some level of analysis in the fused treatment of predicates offered by sophisticated adverbialism that would be able to account for the validity of some inference in the vicinity of (iii) and (iv), as well as for the phenomenon of type-similarity of intentional experiences that we just discussed.

Kriegel's central idea is to treat adverbial modifications as instantiating a determinable-determinate structure. The relation between determinable and determinates is best introduced by means of examples. "Being colored" is a property instantiated by many things. Whatever is colored, however, also instantiates a more determinate property than simply being colored—namely, it is either red, or blue, or green, or whatever color. The property of "being colored," hence, is a determinable property and any determinate color stands in a determinable-determinate relation with that property. Note then that the

determinable-determinate relation goes further down the line. Indeed, one cannot be, say, blue, without instantiating a certain shade of blue, *e.g.*, indigo or cyan.

Let us now see how Kriegel thinks we can apply this distinction to adverbialism. Kriegel bases his rejoinder on some inferences that we regard as valid, such as the following:

(vii) I am eating a strawberry.
(viii) Hence, I am eating a berry.

The inference that goes from (v) to (vi) appears to be valid but it is not certainly so on the basis of the claim that "berry" is a syntactic component of "strawberry." Rather, the claim is that the inference runs through a relation that stands between the concepts "strawberry" and "berry," such that anything that is a strawberry also is a berry. Accordingly, the above inference can be regarded as an enthymematic version of the following:

(ix) I am eating a strawberry;
(x) Everything that is a strawberry is a berry;
(xi) Hence, I am eating a berry.

As Kriegel puts it,

> To me, it seems that what makes the last inference acceptable is some tacit principle to the effect that the property of being a berry is a genus, or determinable, of which the property of being a strawberry is a species, or a determinate. So the full reasoning is something like this: I am kicking a strawberry; whatever is a strawberry is a berry; therefore, I am kicking a berry. (Kriegel 2011: 162-163)

The idea is that the above inference from (v) to (vi) that was treated as invalid can now be treated as valid once it is understood as an instance of the same kind of determinable-determinate structure used to go from (ix) to (xi). That is, "thinking-green-dragon-wise" can be regarded as a determinate of the determinable "thinking-dragon-wise" and, hence, the inference can validly go through, as follows:

(vii) I am thinking green-dragon-wise;
(viii) Everything that is an instance of thinking green-dragon-wise is an instance of thinking dragon-wise;
(ix) Hence, I am thinking dragon-wise.

This very same idea can now also be applied to answer Moran's objection. "Thinking-Obama-wise" and "Thinking-Queen-Mother-wise" can be said to share something that, *e.g.*, "perceiving-Obama-wise" cannot share: they are both determinates of the determinable "thinking."

Kriegel's rejoinder seems to be quite powerful as it is able to provide an answer to both Jackson's and Moran's objection. Is it sound, however? As we shall see in the next section, there are good reasons to doubt that it is.

7. Objections to Adverbialism, Round 2

In this section, we shall discuss two objections to adverbialism. The first refers to some consequences of the endorsement of the notion of phenomenal foreignness, while the second one targets Kriegel's application of the determinable-determinate distinction to adverbialism.

7.1 Adverbialism and the Ganzfeld Effect

The first objection we shall mention targets the claim that phenomenal foreignness is sufficient to account for aboutness. It relies on the so-called Ganzfeld effect, *i.e.*, the experience that one undergoes when one is presented with a completely uniform visual field like, for instance, when one is stuck in a very thick fog. The experience can also simply be induced artificially by filling a room with smoke by means of a smoke machine.

The description of the Ganzfeld effect is of course open to discussion, but the one that seems to be the most adequate is something like the following. When experiencing the Ganzfeld effect, one is uniformly appeared in a white way, without any sense of depth. That is, one might be said to experience pure immanent whiteness, without this whiteness appearing to be a characteristic of anything. Moreover, in a way that is much less controversial than the way in which one is appeared as if everything were black when one closes one's eyes, this experience of pure whiteness is certainly experienced as being foreign. But shall we *also* say that such an experience is about something? It seems quite reasonable to deny that this is the case, and if this is correct, then it seems that the fact that an experience instantiates an adverbial modification such as foreignness is sufficient for it to be intentional.

This formulation of the objection presupposes that there are appearances that are not intentional, and one might regard that as problematic. Note, however, that this objection is consistent with everything that we have said so far in the previous chapters. Indeed, in Chapter 1, we said that for one to be the bearer of an intentional state is for one to be appeared in some way. That said, we have not claimed that intentionality exhausts the realm of appearances.

7.2 Adverbialism and the Determinable-Determinate Relation

In an unpublished paper, Alex Grzankowski (Grzankowski 2017), drawing on independent work on adverbialism by Funkhouser (Funkhouser 2006), suggests that Kriegel's solution does not work. Granzkowski's argument rests on two truisms about the determinable-determinate relation upon which Kriegel's rejoinder infringes.

The relation between determinable and determinates has two directions: one that goes from the determinable to the determinate, and one that goes from the determinate to the determinable. Each direction is governed by a truism. Let us begin with the upward direction, from determinate to determinable.

T1: An object instantiating a determinate also necessarily instantiates every determinable that that determinate falls under.

Take the following example: "That piece of cloth is cyan." "Cyan" is a determinate of blue. Hence, by instantiating the determinate "cyan," "that piece of cloth" also instantiates the property of being blue.

T1 is not problematic for adverbialism. In fact, T1 is the very principle that Kriegel appeals to when arguing that adverbialism can account for the above inference from (v) to (vi). However, let us now look at the principle that rules the downward direction of the relation, the one that goes from the determinable to the determinate, namely:

T2: An object instantiating a determinable must also instantiate some determinate under that determinable.

Take the following example: "That piece of cloth is blue." Something that is blue, however, cannot simply be blue but must instantiate some

determinate shade of blue. Hence, whatever shade of blue "that piece of cloth is," it must instantiate some determinate of the determinable "blue."

T2 is problematic for adverbialism. Take the case of thinking about a green dragon. According to sophisticated adverbialism, this instance of intentionality must be paraphrased as follows: "I am thinking green-dragon-wise." "Thinking green-dragon-wise" is a determinate of "thinking dragon-wise." But it also is a determinable of its own, as sophisticated adverbialism would paraphrase thinking about a big green dragon as follows: "I am thinking big-green-dragon-wise." Now, if T2 is true, then one cannot certainly think green-dragon-wise without also thinking some determinate, such as thinking big-green-dragon-wise. But this only unnecessarily multiplies thoughts, as it is certainly possible for someone to think about a green dragon without thereby thinking about a big green dragon.

In the light of these observations, Kriegel's application of the determinable-determinate relation to adverbialism then faces the following dilemma: either there are good, non *ad hoc* reasons to consider that T2 applies only to a sub-kind of determinables and the determinable-determinate relation can be applied successfully to adverbialism; or there are no such reasons and the application of the determinable-determinate relation to adverbialism is problematic. As Kriegel does not provide any reasons in favor of the first disjunct, and as the prospects for its truth seem dim, one might be allowed, if only in a *prima facie* way, to endorse the truth of the second disjunct. To conclude, then, adverbialism does not constitute an available option for the proponent of a monadic theory of intentionality.

6. Conclusion

In this chapter, we discussed a first version of a theory of intentionality that denies that the underlying metaphysical nature of intentionality is that of a relation, namely adverbialism. We saw that adverbialism as applied to a theory of intentionality proceeds in two steps. First of all, it proposes to systematically translate intentional sentences into adverbialist paraphrases. Second, it proposes to account for the truth of these paraphrases in terms of adverbialist truth-makers, *i.e.*, conscious mental states, or experiences, that instantiate adverbial modifications.

Accordingly, an experience being intentional is cashed out adverbially in terms of an adverbial modification of being intentionally directed *F*-wise, where the intentional nature of this modification is understood in terms of the further notion of phenomenal foreignness and, when required, particularity.

We discussed two series of objections against adverbialism. The first round consisted of two classical objections, one by Frank Jackson and one by Dermot Moran. As we saw, Kriegel tries to answer these objections by cashing out adverbialism in terms of a determinable-determinate structure. We raised two further objections against this second version. The first is that it saddles intentional experiences with much more thoughts than is needed. The second is that the notion of phenomenal foreignness is not sufficient to account for intentionality.

Chapter 6: Polyadism

> The claim that there are non-existent intentional objects is not an ontological claim. It is simply another way of saying that we can genuinely think about things that don't exist, and that we can think about them in the same way in which we think about existing things.
>
> (Crane 2013: 5)

1. Introduction

So far, we have discussed three possible interpretations of the FIS: the IOM, the CM, and adverbialism. Each of them raised some concerns. What is the next step? In this last chapter, we discuss a recent new addition to models of intentionality, proposed by Tim Crane in his recent *The Objects of Thought* (Crane 2013). As we shall see, Crane tries to merge the best of two worlds: the descriptive power of the IOM and the theoretical economy of the CM and adverbialism.

Indeed, despite their differences, the CM and adverbialism share a same nature that sets them apart from the IOM. Namely, they are representational theories of intentionality. In other words, each of these two models assumes that to be the bearer of an intentional experience is not to stand in relation to that upon which one is

intentionally directed, but rather to be the bearer of a content that may or may not match the world. As such, being the bearer of an intentional experience understood in such a representational way does not require us to provide intentional experiences with something, either existent or not, such that they can stand in the "being about" relation to that thing.

The IOM, on the other hand, is a non-representational theory. Indeed, according to the IOM, there is no possibility of having an intentional state that does not match a mind-independent something that is somehow out there. All intentional experiences are directed upon something that features in the overall inventory of what there is, although some intentional experiences are about things that exist while some others are about things that do not exist.

The property of being representational alleviates many of the troubles faced by the IOM as it frees the theory of intentionality from the burden of endorsing Meinongianism. But, as we saw, being representational is not in itself a guarantee of being an adequate model of intentionality. That said, one might claim that representational theories are intrinsically better than non-representational ones. Intuitively, they seem to get something basic intuitively right: Thinking about something that does not exist does not amount to one's mind latching onto something that does not exist but is somehow still there, in a mind-independent way. The IOM, with its endorsement of Meinongianism, goes astray on that very point.

Standard forms of representationalism like the CM and adverbialism, however, are committed to a supplementary claim. That is, they do not only claim that intentional experiences are representational; they also claim that aboutness is a notion that is subject to success and failure. Indeed, even though these theories propose to carve a notion of aboutness that is generally instantiated

by all intentional experiences, the non-technical sense of being about emphasized in Chapter 4 remains instantiated only by a sub-kind of intentional experiences. In that sense, according to such theories, an intentional experience that is about something that does not exist fails to be about anything.

This supplementary claim, however, is hard to reconcile with an adequate description of intentional experiences. Indeed, all intentional experiences are about something in the sense of being directed upon something, or having a main topic—in other words, all intentional experiences have an intentional object. This is simply a basic phenomenological fact about intentional experiences. That is, from a phenomenological point of view, it seems utterly mysterious to claim that an intentional experience may fail to be about something, in the intuitive sense of being about. As Tim Crane puts it,

> Our mental life seems to involve the presence to the mind—or apparent presence to the mind—of things in the world. These things can be mental or material, concrete or abstract, and—so I say—existent and non-existent. I believe that there are general characteristics of intentionality which apply to all, or almost all, intentional states and episodes. One of these is that every intentional state or episode has an object—something it is about or directed on. (Crane 2013: 4)

The IOM is right on that point: Aboutness is not a notion subject to success or failure. It is, however, mistaken on the claim that intentionality is not representational. The distinction between these two points, a theory of intentionality being representational and a notion of aboutness being subject to success and failure, makes room for the following two questions:

(a) Is a literal interpretation of the FIS bound to be non-representational, like the IOM?; and

(b) And if not, can we develop a modified version of the IOM that treats intentionality as representational without treating aboutness as a notion subject to success and failure?

In the recent work mentioned above, Tim Crane argues that we can negatively answer question (a) and positively answer question (b). The resulting model of intentionality is what we shall call, for reasons that will become clear below, polyadism.

The plan of the chapter is the following. In the next section, we come back to the distinction between a non-representational theory of intentionality and making use of a notion of intentional directedness that is not subject to success or failure. We argue that these two notions can be unraveled and, hence, that there is room for a theory that, like Crane's, endorses the latter feature without endorsing the former.

In section 3, we introduce and explain Crane's metaphysical distinction between two kinds of properties and relations: first, substantial properties and relations; and, second, non-substantial properties and relations. As he sees the matter, the constitutive condition for an experience to be intentional is merely that it be a *non*-substantial relation. For a non-substantial relation to obtain, however, it is not necessary that there be something to which one stands in relation. As such, then, intentional experiences that are about things that do not exist can have a somewhat relational nature without thereby being relations to non-existent objects.

In section 4, we subsequently distinguish between two ways in which a proposition can be literally true, and explain how this distinction can account for the fact that sentences such as "John

thinks about Pegasus" can be literally true without there being an non-existent Pegasus to which John stands in relation. In section 5, we then turn to how Crane's model is supposed to account for the six main features of intentionality that we identified in Chapter 1 before, in section 6, discussing some problems of polyadism.

2. Representation, Success, and Intentionality

According to the IOM, all instances of intentionality are instances of the following schema:

> **IOM-Schema:** "*x* is about *y*" amounts to "*x* stands in the relation of 'being about' with an intentional object *y,* and for some content *C*, *C* presents *y* to *x*."

As we saw, the generalized instantiation of this schema by each instance of intentionality, including instances that are about things that do not exist, requires the endorsement of Meinongianism, according to which there are mind-independent non-existent intentional objects to which we can stand in genuine relations.

As such, the IOM constitutes a *non-representational* theory of intentionality, in the sense that for each instance of the above schema, there is an element of the overall inventory of what there is that can serve as the relata of the "being about" relation. Take the case of John thinking about Pegasus. According to the IOM, for John to think about Pegasus is for him to really stand in a relation to Pegasus, as presented under some content *C*. In other words, Pegasus itself serves as a ground of the truth of this specific instance of the above schema. The IOM, in sum, constitutes a non-representational theory of intentionality because all intentional experiences put us in a

relation—a *genuine* relation—with a mind-independent element of the overall inventory of what there is.

Hence, what is characteristic of the IOM is that some facts about the mental realm, *i.e.*, that we can think about things that do not exist, are explained by virtue of facts about the non-mental realm, *i.e.*, that there are, in a way that is independent of the mental, more things in the overall inventory of what there is than those that exist and to which we can stand in relation.

It will become essential to what comes next to distinguish between this first feature of the IOM, being non-representational, and an important corollary of this claim. Indeed, some might find it too obvious to be emphasized, but the IOM is not only a non-representational theory of intentionality; it also one in which being intentionally directed upon something is not, unlike in the CM or in adverbialism, something that is subject to success or failure. That is, according to the IOM, whenever one is undergoing an intentional experience, then there is something upon which one is intentionally directed: an intentional object. As we saw, this is not the case for the CM and adverbialism, according to which only referentially successful intentional experiences possess this feature.

That said, one might be willing to ask the following question: What is the relation between the notion that intentionality is non-representational, and the one that being intentionally directed upon something is not subject to success or failure? If we look at the relation that goes from being non-representational in the sense defined above to not being submitted to success or failure, then it seems quite clear that the former entails the latter. In the IOM, there is no room for success or failure of intentional directedness because the theory guarantees a mind-independent intentional object for each intentional experience. But what about the converse relation? Does

the claim that being intentionally directed upon something is not subject to success or failure also entail non-representationalism?

One might be forgiven for believing that the truth of this second implication is as trivial as the truth of the first one. Indeed, how could an account of intentional directedness not subject to success and failure be cashed out independently of the IOM? In his recent book *The Objects of Thought*, however, Tim Crane drives a wedge between these two claims (Crane 2013). In a nutshell, he develops a model of intentionality that both makes an irreducible use of the notion of intentional objects *and* rejects the IOM's non-representationalism. In that sense, his model constitutes a middle ground between the representational theories that we have already discussed—*i.e.*, the CM and adverbialism—and the IOM.

Before addressing the details of Crane's proposal, let us try to convey the general spirit of his position. At the core of the problem of intentionality lie sentences or propositions that feature an existent intentional subject, an intentional verb, and a non-existent intentional object, such as "John thinks about Pegasus" or "John fears Beelzebub." Indeed, according to a theory of intentionality true to the phenomenological facts, such sentences would have to be literally true. But their literal truth, or so it seems, commits us to philosophical monsters such as Meinongianism. Phenomenological sanity therefore seems to come at the price of theoretical insanity, and, justifiably, we have a tendency to dump the former in order to save ourselves from the latter.

In other words, according to Crane, at the core of the problem of intentionality lies the following question:

(Q) How can a predication, either monadic or polyadic, formed by means of at least one singular term that lacks an existent referent be literally true?

Note that Crane's account takes a slight shift from former theories of intentionality. Indeed, he does not start with the assumption that, in the IOM, intentionality is a relation to intentional objects. Nor, alternatively, does he begin with the idea that a correct interpretation of the FIS cannot be a literal one, like the CM or adverbialism. On the contrary, he starts with the claim that it is a bedrock fact that there are literal truths about non-existent objects, and that it is part of the task of a theory of intentionality to provide us with an account of that fact. As such, then, he genuinely believes that the problem of non-existence as faced by the IOM must be solved, and, unlike in the CM and or adverbialism, merely be dissolved.

Moreover, according to Crane, we take (Q) to be philosophically problematic—so problematic, in fact, that we sometimes prefer to resort to extreme means such as denying that a sentence like, "John thinks about Pegasus," can be literally true—because we believe that the following two claims are true:

> **Truth 1:** If a sentence of the form "*a* is *F*" is literally true, then the ground for this truth must be that there is some object, *a*, that instantiates *F* (and something similar holds for *polyadic* predications); and
> **Truth 2:** If *a* does not exist, then this ground for truth mentioned in Truth 1 must be partly constituted by a non-existent object, *i.e.*, *a*.

The standard dialectic as to why a sentence like "John thinks about Pegasus" turns out to be problematic somehow goes as follows: We want to stick to Truth 1 but cannot cope with Truth 2. But since our problem with Truth 2 merely follows from the temptation to read sentences such as, "John thinks about Pegasus," in a literal way, then

we simply dump the claim that, "John thinks about Pegasus," can be literally true and take on board a theory that makes it, at best, *non-literally* true. We are then free to keep Truth 1 and we can stop bothering about Truth 2.

From Crane's perspective, however, this dialectic is problematic because, despite the advantage of the move just described, the intuition—or, better, the phenomenological fact—that sentences like, "John thinks about Pegasus," can be literally true simply does not go away by sticking to Truth 1 and dumping Truth 2. And, decisively, we will not be able to put the theory of intentionality to rest—if ever—without having found the means to pay dues to the literal truth of at least some of such sentences.

Crane, however, proposes an alternative diagnosis of the situation. Namely, if we want to recognize that some simple sentences formed by means of singular terms that are about things that do not exist can be literally true, then we are stuck with Truth 2 *only* if we unrestrictedly stick to Truth 1. If we were somewhat able to account for the claim that a sentence like "*a* is *F*" can be literally true without this claim being theoretically tied to Truth 1, then we would able to bypass the problem of having to face Truth 2.

The resulting theory, moreover, would count as representational with respect to the way in which we characterized this notion above. Indeed, if a theory of intentionality is non-representational just in case the literal truth of a sentence like, "John thinks about Pegasus," requires that Truth 1 and Truth 2 hold, then Crane's model turns out to be a representational one.

What, however, would be required for us to be able to let Truth 1 go? According to Crane, three essential building blocks are required, namely the following three theses:

(T13) A theory of quantification is not committed to a classical logic reading of quantifiers;

(T14) For a restricted range of non-existence entailing properties, Meinong's Principle of Independence of *Sosein* from *Sein* holds; and

(T15) Some sentences of the form "*a* is *F*" can be literally true, without there being any ground of truth for such sentences to the form, "*a* that is *F*."

If a sentence like, "John thinks about Pegasus," is literally true, then it follows that some non-existent objects can be thought about. Hence, we need a theory of quantification that allows us to quantify over non-existents. Such an account, as we saw in Chapter 3, is available, and in this chapter we shall simply presuppose this previous discussion and focus instead on (T14) and (T15).

Before turning to the case in favor of them, let us try to see how a theory of intentionality that takes on board (T13)-(T15) would look. On the predicative side, Crane proposes that we distinguish between two kinds of properties: substantial and non-substantial ones. A restricted class of non-substantial properties, representation-dependent ones, are not existence entailing and can be used to make a restricted version of the Principle of Independence true. And since being thought about in some way counts as a representation-dependent property, then non-existent intentional objects can be thought about in a way that is indifferent to their ontological status.

On the alethic side, Crane distinguishes between two kinds of truths: namely the truth of propositions that result from the predication of substantial properties, and the truth of those that result form the predication of non-substantial properties. According to him, only the first kind must be accounted for in terms of Truth 1. True predications of non-substantial properties can be accounted for

by grounds of truth that have, to use Crane's own metaphorical term, a "shape" that is different from the truth that they ground (Crane 2013: 65). In other words, "John thinks about Pegasus," can be literally true, according to Crane, without there being a ground of truth to the shape, "Thinks about (John, Pegasus)," even though there *is* a proposition with that shape that is literally true.

Moreover, this overall strategy counts as representational because, contrary to the IOM, it is not the case that if, say, "John thinks about Pegasus," is true, then this is made true by there being a mind-independent non-existent intentional object to which John stands in relation. All of the grounds exist and are entirely constituted of things that exist. According to Crane, it is one of these that, in a way that we shall discuss below, accounts for the literal truth of, "John thinks about Pegasus."

Aiming to summarizing his overall strategy in general terms, Crane writes:

> It is natural to distinguish between the truth of some sentence or proposition, and what we might call picturesquely (but I hope quite intelligibly) the 'shape' of the reality described by that truth. A true sentence might straightforwardly mention certain kinds of thing—a certain particular object, a fact, proposition, or property or some other kind of thing. So it might seem as if, in maintaining the truth of this sentence, we are committed to such entities. But such straightforward commitment can be maintained together with a different view about the shape of the reality described by that truth. It may seem at first sight that we are committed to propositions or facts because of the sentences we hold true, but the underlying reality need not match this commitment in any simple or straightforward way. (Crane 2013: 65)

Leaving the details aside for the time being, such an account would, first of all, count as making use of a notion of being intentionally directed upon something that would not be subject to success or failure, as, *e.g.*, if John thinks about Pegasus, then, "John thinks about Pegasus," would be literally true. It would, furthermore, also count as a representational theory of intentionality insofar as for John to think about Pegasus would not amount to John standing in relation to a mind-independent non-existent object, but rather to John instantiating a complex representational structure. Having sketched the great lines of Crane's view, let us now turn to a more detailed discussion.

3. Substantial *vs.* Non-Substantial Properties

In this section, we tackle the issue of (T14) and discuss Crane's introduction of the distinction between substantial and non-substantial properties. As we shall see, the matter can only be tackled somewhat indirectly, through some remarks on kinds of properties and on the conditions under which the introduction of a new distinction between kinds of properties is justified.

Theories of properties generally specify not only that there are properties, but also that there are different *kinds* of properties. Some standardly recognized kinds of properties are logical properties, existential properties, categorical properties, dispositional properties, physical properties, mental properties, sortal properties, instantiated properties, non-instantiated properties, and so on. Of course, not all theories of properties recognize such a great variety of properties, but they do generally recognize *some* varieties of properties.

But who decides which kinds of properties there are, and on what grounds? Whoever introduced the notion of property into philosophy did not at the same time introduce all kinds of properties.

Rather, over time philosophers have come to discover distinctions between different kinds of properties, and have thus enlarged the list of the candidate kinds of properties. Hence, no one decides which kinds of properties there are, or should be, besides the ongoing process of philosophical theorizing.

This, however, raises a question. Under which conditions is one allowed to introduce a new kind of property? This is a difficult question to answer, but we might apply a rule of thumb that might be said to apply generally to the introduction of any kind of philosophical distinction: One is allowed to introduce a new distinction when it allows for a real philosophical difference, *i.e.*, when it allows us to say true, or at least justified, substantial things that we would not be able to say without that distinction. Let us illustrate this claim by means of an example that pertains specifically to the introduction of a new distinction between kinds of properties.

3.1 Lewis on Natural and Non-Natural Properties

David Lewis is well known for holding a specific version of class-nominalism, *i.e.*, the doctrine that properties can be identified with classes of things. Class-nominalism, however, is beset by a simple objection that follows from the extensionality of classes: If two classes possess exactly the same members, then they are the same class. Take, for instance, the class of the creatures that have a heart and the class of the creatures that have kidneys. As it turns out, these two classes have exactly the same members and, as a result, having a heart and having a kidney turns out to be exactly the same property —a sure proof that class-nominalism, at least in this naive version, is mistaken.

Lewis, however, defends a more sophisticated version of class-nominalism. Indeed, he claims that the classes with which we shall

identify properties are trans-world classes of concrete individuals located at different, causally isolated, concretely existing possible worlds, in accordance with his infamous doctrine of "modal realism" (Lewis 1973, 1986). That is, classes constituted not only of individuals located in our worlds, but also of numerically distinct individuals located in other possible worlds. Classes of individuals thus conceived then allow Lewis to dismantle the co-extensionality objection as being based on an unjustified restricted interpretation of the universal quantifier. Indeed, the universally quantified sentence, "All creatures that have a heart also have a kidney," is certainly true at *our* world, but it turns out to be false when "all" is allowed to be interpreted unrestrictedly. When unrestricted, "all" ranges over the totality of possible worlds and it is simply false that in every possible world any creature that has a heart also has a kidney, as there is no metaphysical or logical impossibility in there being a possible living creature that, for instance, possesses a heart but no kidney. Hence, if the class of the creatures that have a heart is co-extensional to the one of the creatures that have a kidney *at our world*, the charge of co-extensionality collapses once we consider their corresponding trans-world classes.

However, if the machinery of modal realism allows Lewis to avoid the objection of co-extensionality, he must still face another objection against class-nominalism, which we may call the objection from over-proliferation. Indeed, even without taking on board the infamous axiom of unrestricted comprehension, set theory still allows for the formation of arbitrary classes. This, of course, is not a problem in itself, but if one is willing to use a theory of classes to account for a nominalist theory of properties, as Lewis does, then this set-theoretic generosity will cause an over-proliferation of properties, since to each class, independently of any consideration with respect to its arbitrariness, will correspond a property. As a

result, one will end up with so many properties that one might wonder whether properties could still be said to play the substantial philosophical role that they are supposed to play.

Indeed, properties are notably said to be necessary to account for natural resemblances and similarities, *i.e.*, two spheres are identical in kind because they share the same property of being spherical. But if we can form arbitrary classes and if, additionally, to each arbitrary class corresponds a property, then pretty much anything can be said to resemble anything else. My keyboard and the Matterhorn, for instance, can be said to resemble each other because they share the property of being a member of the arbitrary class π {my keyboard, the Matterhorn}. This kind of resemblance, however, has little to do with the one that is shared by two spheres, and the mere idea that properties can account for identity in kind between particulars seems to have been brought to a *reductio*.

Faced with this problem, one could espouse different solutions. For instance, one could simply give up on the identification of properties with classes. Lewis, however, chooses another path. He settles on the identification of properties with classes, but proposes to distinguish between two kinds of classes, and hence two kinds of properties: *natural* classes and properties on the one hand, and *non-natural* classes and properties on the other (Lewis 1983).

Equipped with this distinction, Lewis can then say that two spheres resemble each other in the sense of being members of the natural class of spherical things, while my keyboard and the Matterhorn resemble each other in the sense of being members of the non-natural class π. Natural properties can then be said to play the substantial role generally attributed to properties, namely carving out natural resemblance classes, while non-substantial properties can propagate themselves without the risk of damaging Lewis's overall account.

We do not need to concern ourselves with the specifics of Lewis's proposal here, nor with its evaluation (though, in fairness, it is important to note that Lewis does not think that the distinction between natural and non-natural classes is primitive). What matters for us, rather, is the fact that Lewis's introduction of a distinction between natural and non-natural classes and properties can serve as an example of an introduction of a new distinction between kinds of properties in a way that is, we shall assume, justified. Indeed, the introduction of this distinction allows him to defend his version of class‑nominalism, an important nominalist contender, in the face of the substantial over-multiplication objection mentioned above. As such, the introduction of this distinction allows him to say something that is philosophically substantial that he would not be able to say otherwise.

3.2 Crane on Substantial and Non-Substantial Properties

This nod to Lewis gives us a first key to Crane's treatment of intentionality, and in particular to the above thesis (T14). Namely, Lewis's account can serve as an example of the claim that we can, where it is needed and justified, introduce a new distinction between kinds of properties. Crane himself, however, is not concerned with Lewis's distinction between natural and non-natural properties, but rather with a distinction between what he calls "substantial" and "non-substantial" properties (Crane 2013: 66). He characterizes this distinction as follows:

> [The distinction between substantial and non-substantial properties is] the distinction between properties and relations which can be read simply of the surface structure of the sentences used to describe them, and the properties and relations that cannot, those that

characterize the nature of real existing things. These latter properties are the substantial properties. (Crane 2013: 66)

The general idea is the following. Take as a start a schema like "*a* is *F*." From this schema, we can obtain a syntactic conception of properties by means of generalizing on the second variable, *F*, as follows: A property is anything that can be substituted for *F*. This schema can then be freely extended to polyadic properties.

On the basis of this syntactic conception, we can then distinguish between two kinds of true predicative sentences along the following lines. For each true sentence of the form "*a* is *F*," we can infer the following: "There is an *F* that *a* instantiates." We might, however, distinguish between two kinds of such inferences. On the one hand, there are these inferences that we perform simply on the basis of the fact that the original sentence is true. Accordingly, we do not commit ourselves to a substantial conception of the property *F* in question and do not claim that a substantial empirical or metaphysical inquiry would deliver us the result that there is, in a substantial sense, a property *F*. Rather, we perform this inference on the sole ground that the original simple sentence is true. The predicative sentences on the basis of which we can perform such inferences are the ones that are formed by means of a *non-substantial* property.

On the other hand, there are those inferences that go from a sentence of the form "*a* is *F*," to one of the form, "There is an *F* that *a* instantiates," that we do not endorse simply on the basis of the truth of the original sentence. Rather, we endorse these inferences only after a substantial empirical or metaphysical inquiry, thereby committing us to the idea that such an *F* possesses a substantial nature about which a theoretic discipline can inform us. These are the substantial properties.

Here is an example:

(1) Obama is famous;
(2) Hence, there is a property that Obama possesses, *i.e.*, being famous.

We know that (1) is true. How? We read the newspapers, watch TV, discuss with our peers, and as a result simply know that Obama is famous. From this, we can then infer (2). But does the inference from (1) to (2) commit us to there being a substantial property of being famous that, *qua* property, would be on par with, say, the property of being positively charged? The idea of a non-substantial property is that we do not need to be thus committed to a substantial property of being famous. We might simply say that (1) is a case of a sentence saying something literally true about some subject matter, without thereby committing us to a substantial conception of that something that is said to be literally true of some subject matter. However, this is not to deny that Obama also possesses substantial properties, some of which, we might add, might account for the truth of (1).

According to Crane, there is an important ontological difference between substantial and non-substantial properties: Namely, a true predication formed out of a non-substantial sentence is not ontologically committing. As he puts it, "Ontology, the study of being, is concerned with properties and relations in the substantial sense, not [the non-substantial sense]." (Crane 2013: 67) To recognize as literally true that a certain object possesses a certain non-substantial property does not, hence, amount to making an ontological claim.

We must then make a distinction between those predicative sentences that are ontologically committing and those that are not. The former, presumably, cut the world along its real joints, *i.e.*, they

cut the world along the joints that are discovered by substantial empirical and metaphysical inquiries. But what about the latter? Well, presumably, they only report truths about the way in which the world is experienced or represented. We experience the world as being such that it is literally true that Obama is famous, for instance. A denial of this literal truth would simply amount to an amputation of the way in which we experience the world as being. But in order to accept this claim, we do not need to affirm that the property of being famous is, *qua* property, a non-par with a substantial property.

According to Crane, the introduction of this distinction between substantial and non-substantial properties constitutes a first step in the direction of the truth of the above (T14). The next crucial step is the claim that even though non-existent objects cannot instantiate substantial properties, they can, for at least some sub-kinds of them, instantiate non-substantial properties. As he puts it,

> Equipped with these distinctions, we can say that non-existent objects can have properties and stand in relations in the pleonastic, non-substantial sense. Something, *'F'*, can be true of a non-existent object, and in the pleonastic sense this object has that the property *F*. So I reject the claim that (in Malebranche's words) 'nothingness has no properties'. (Crane 2013: 67-68)

In Chapter 3, we discussed Crane's objection to Meinongianism based on the idea that the Characterization Principle (CP) was false because many properties are existence entailing ones and therefore cannot be instantiated by a non-existent object. Let us recall some details of the objection. First, the CP is the principle that for each set of nuclear properties, there is an object that corresponds to that set. According to the CP, then, "Pegasus is a horse," is true, granted that Pegasus is characterized as a horse. According to Crane, however, this claim is

inconsistent with the one that there are some plausible requirements on being a horse that a non-existent object will never be able to fulfill. As such, then, Crane argues that the CP must be rejected. But does this mean that the mere idea that non-existent things can be the objects of true predication must be rejected?

According to Crane, the distinction between substantial and non-substantial properties allows us to see that even though the CP is false, there is no reason to deny that non-existent objects such as Pegasus can still instantiate a restricted range of properties, namely non-substantial ones. Indeed, if, "Pegasus is a horse," is true, since Pegasus does not exist, "Pegasus is famous," might still be true.

Crane distinguishes between three kinds of sentences about non-existents that can be true, and hence between three different kinds of properties that can be instantiated by non-existent objects. The first kind of sentences is negative existential ones, such as, "Pegasus does not exist." According to Crane, such a sentence is true, and its truth simply follows from the fact that, "Pegasus exists," is false, following the doctrine that the negation of a falsity is a truth. The second kind of sentences is those of trivial identity, such as, "Pegasus is Pegasus," which, he claims, simply follows from the logical truth that for all x, $x = x$. Finally, the third kind is predicative sentences formed from what he calls, following Colin McGinn (McGinn 2001), "representation-dependent properties." (Crane 2013: 68) This last kind of sentence forms the great bulk of truths about non-existents objects.

A representation-dependent property is a property that depends on the fact that an object is represented in some way. In the above (1), "Being famous," obviously, is such a property. Indeed, that someone is famous amounts to nothing but the fact that this person is intensely thought or spoken about. Withdraw representation from the world and you thereby withdraw the property of being famous.

The strategy is then the following. We know that at least some sentences about non-existent objects are true. But by distinguishing between substantial and non-substantial properties, first of all, and by almost completely restricting properties instantiated by non-existent objects to non-substantial ones, we are able to bypass the objection from nature discussed in Chapter 3 against the claim that non-existent objects can instantiate properties. Indeed, the thought is that something does not need to exist to be represented somehow. For instance, John can think about Pegasus independently of Pegasus's ontological status.

As such, then, the introduction of a distinction between substantial and non-substantial properties allows us to make a case in favor of (T14). As such, we can account for the claim that non-existent objects can instantiate at least some properties and relations, namely non-substantial ones.

3.3 Aboutness as a Non-Substantial Relation

Traditional theories of intentionality distinguish between two options for a theory of intentionality. Either a theory is *relational*, like the IOM or the CM, or it is *monadic*, like adverbialism. According to Crane, however, this disjunction is not exhaustive. A third option that is possible with the endorsement of the distinction between substantial and non-substantial properties.

Indeed, if we can distinguish between substantial and non-substantial monadic properties, we can certainly also distinguish between substantial and non-substantial relations. Hence, the denial that intentionality is a relation can actually mean either one of two things: either that intentionality is a monadic property, or that intentionality is a non-substantial relation. Crane argues that it is the later. As he puts it,

> …there are many kinds of relations, but some are substantial and some are not. A non-substantial relation might simply consist in the truth of a relational predication. There's no reason to deny, for example, that whenever a relational proposition 'aRb' is true, the relation R holds between *a* and *b*. So when I said in chapter 1 that intentionality is not a real relation, what I meant was that it is not a substantial relation. (Crane 2013: 66)

This means the following. Intentionality is a non-substantial relation, or a mere polyadic predication. Sentences describing intentional experiences, such as, "John thinks about Pegasus," can be literally true and can be taken to express propositions with the form "Thinking about (x,y)." However, the fact that such a polyadic predication obtains is *not* equivalent to the claim that a relation obtains. In other words, a sentence like, "John is thinking about Pegasus," can be literally true in the sense of being analyzed as "Think about(John, Pegasus)," but this is not equivalent to the claim that a relation obtains between John and Pegasus.

The distinction between substantial and non-substantial properties is supposed to allow a sentence like, "John is thinking about Pegasus," to be literally true. But how can such a sentence be true if no relation obtains between John and Pegasus? According to Crane, this is where we should let go of our hold on Truth 1. This, however, will require a defense of the above (T15), which will be discussed in the next section.

4. Truth About Non-Existents

The dialectic since the beginning of this work can be summarized as follows. We know that some sentences like, "John thinks about

Pegasus," are true, even though Pegasus does not exist. The question is then how such sentences can be true in light of this fact about Pegasus's ontological status. The culprit, as we discussed at length, lies in the interpretation of the little word "about." The IOM reads it literally as being a relation between an intentional experience and an intentional object, whereas the CM and adverbialism read it non-literally, each in their own way.

Despite their differences, each of these theories can nonetheless all be regarded as sharing a common assumption, namely that the thesis Truth 1 already introduced above is true, namely:

> **Truth 1:** If a sentence of the form "*a is F*" is literally true, then the ground for this truth must be that there is some object, *a*, that instantiates *F* (and something similar holds for *polyadic* predications).

This endorsement leads the IOM to embrace Meinongianism and leads the CM and adverbialism in a rather different direction, *i.e.*, a paraphrastic treatment of the FIS with the aim of avoiding treating intentional sentences as predicative ones. As we saw in the previous section, Crane aims to offer a third way, *i.e.*, a literal interpretation that remains very close to the IOM without presupposing that intentionality is a relation. Indeed, he distinguishes between substantial and non-substantial properties and relations, and argues that intentionality is an instance of the latter.

A non-substantial relation is not a monadic property insofar as it remains syntactically polyadic. It is also not a relation, in the sense that it does not require a relational fact as a ground of truth. Well, good enough, one might say, but what does this exactly *mean*? Our interpretative strategy shall be to take Crane's strategy as having two parts: one part about *meaning* and one part about *truth*.

The part about meaning is the following. The metaphysical work achieved by the distinction between substantial and non-substantial properties offers new opportunities for the analysis of intentional sentences, *i.e.*, we no longer need to choose between analyzing a sentence like, "John thinks about Pegasus," either as a relation of "being about" along the lines of the IOM, or under the form of a paraphrase, like the CM or adverbialism. The game is now a three-way one: either it is a relation, or a monadic property, or a non-substantial relation. That is, we can now say that a sentence like, "John thinks about Pegasus," expresses a proposition of the form "Thinks about (John, Pegasus)," without thereby meaning that intentional experiences are relations to intentional objects.

Second, the part about truth is that not all sentences that are literally true need to have a ground of truth with exactly the same shape as the proposition that they express. For instance, a predicative sentence of the form "*a* is *F*" expressing a sentence of the form "*F(a)*" does not necessarily need to be made true by a ground of truth with the shape, "There is an object *a* that instantiates the property *F*." Indeed, recall the passage already quoted above:

> It is natural to distinguish between the truth of some sentence or proposition, and what we might call picturesquely (but I hope quite intelligibly) the 'shape' of the reality described by that truth. (Crane 2013: 65)

As introduced in Chapter 3 in our discussion of predication, the notion of literal truth is the following: A sentence with a predicative grammatical form can be said to be literally true if, and only if, (i) it expresses a predicative proposition and (ii) this proposition is true. Note, however, that we did not explicitly build more than that into the notion of literal truth. In particular, we did not build any idea

about the *grounds* for such literal truths. Accordingly, we can distinguish between two ideas: first, that a natural language sentence *S* can have a literal interpretation; and second, that this literal interpretation is true by virtue of a principle such as Truth 1.

Crane's claim about truth must then be that a sentence of the form "*a* is *F*" can be literally true even if there is no ground for this truth with a similar shape as that of the proposition expressed by this sentence. We must then reconfigure the notion of literal truth so that some sentences can be literally true without necessarily fitting Truth 1. Crane's fundamental idea in that area is that this alternative notion of literal truth should appeal to the idea of a possible "reductive explanation" of truths (Crane 2013: 124).

The idea is the following. It is common ground in philosophy that we want some sentences to be true without, as Crane puts it, "wanting to commit ourselves to the entities which these claims apparently talk about." (Crane 2013: 122) The standard way to do this is to appeal to a paraphrase strategy. But, he claims, this is not the only available strategy. One can also try to provide a reductive explanation of a truth in a way that does not appeal to the entities that we have identified as problematic. This can be achieved by explaining, or "making intelligible," (Crane 2013: 122) how certain truths about something ontologically problematic can be true by appealing only to non-problematic ontological facts.

What is the difference between a paraphrase and a reductive explanation, however? We have already discussed the question of paraphrases several times. Take, once again, sentences like, "I did it for John's sake," or, "John wore a broad smile." Let us assume that these sentences are true. That said, we are not willing to recognize that there are entities like sakes or smiles to which we can stand in relation. The paraphrase strategy then consists of claiming that these sentences are true, but only non-literally. That is, they are true in

virtue of a paraphrase being true. This paraphrase, moreover, can either be conservative or revisionary, but assuming a descriptive picture according to which they must be conservative, then such paraphrases must then be evaluated in terms of native speakers' linguistic intuitions. A good conservative paraphrase can then be taken to capture the meaning of the original sentence that it paraphrases.

As such, then, the paraphrase strategy, understood as a descriptive strategy, is committed to a central claim about *meaning*. The paraphrase is taken to capture the meaning of the original sentence and, since it is not committed to any entity to which we will not want to be committed, we can conclude that the original sentence is not, in fact, committed to anything problematic either. In other words, a paraphrase strategy deflates our ontological problems by means of fixing the meaning of a problematic sentence. As Crane puts it,

> What we do in these cases is to provide a paraphrase or translation of the original claim, to show that we can mean *what we mean* by it without committing ourselves to sakes [or smiles]. (Crane 2013: 122– italics added)

The major difference between a paraphrase and a reductive explanation as conceived by Crane is that the latter strategy does not start by trying to fix the meaning of a problematic sentence in order to make some ontological problem disappear. Rather, it embraces the literal meaning of a sentence, and tries to ease its potential ontological problem in two steps: First, where possible, we turn substantial predications into non-substantial ones; and second, we find a reductive ground of truth for this sentence—that is, a ground of truth that is unproblematic and from which we "make intelligible"

(Crane 2013: 122) why the original sentence is literally true. In sum, if one tries to give a reductive explanation of a problematic sentence, one is not saying that there is a problem with the literal meaning of that sentence and therefore with it being literally true. Rather, one is saying that a problematic sentence can have its literal meaning and be literally true without thereby being committed to problematic entities that make it true. Crane explains it as follows:

> …in other (more interesting) kinds of cases, we do not provide translations, but rather attempt to describe the world in a way that makes intelligible why the claim is true, while not committing us to entities which we are talking about. (Crane 2013: 122)

The notion of reductive explanation is introduced by means of examples. Crane's main example relies on a sentence like the following:

> (3) Our society puts too much value on making money for its own sake, and it should instead emphasize the value of learning, culture, sport, and good health.

Let us admit that (3) is true. Not everyone, however, is a realist about societies. Methodological individualists, for instance, are not. It is not the place and time, however, to ponder methodological individualism and for the present issue, the following rough and ready characterization will suffice: It is, in a nutshell, the idea that any feature of a society can ultimately be explained in terms of the actions and intentional states of atomized agents. In Margaret Thatcher's famous words, "[a]nd, you know, there is no such thing as society. There are individual men and women, and there are families." (Quoted in Crane 2013: 122)

The truth of (3) creates a problem for methodological individualism. Indeed, are methodological individualists forced to claim that (3) is only non-literally true, *i.e.*, that it is true only because a paraphrase—presumably either a big disjunction of names of people putting too much emphasis on making money or, alternatively, an existentially very complex and long statement—is true?

According to Crane, this is too strong. There is no need for a proponent of methodological individualism to claim that (3) is true only because a paraphrase from which the word "society" has disappeared is true. Indeed, she could well stick to the thesis that (3) is literally true, and not true in virtue of a paraphrase being true, and explain why it is true in a way that does not presuppose more than atomized individuals, their actions, and their intentional states. As Crane puts it,

> What I want to emphasize here is that it is perfectly coherent to say that when someone insists on the truth of the claim 'society places too much value on making money', they do not have to be construed as *saying something that means* that people in society tend to place to much value on making money. Even if this latter claim is a consequence of what they say, this does not imply that there is any way of translating claims about societies into claims about people. Rather, it's that the claim about people *explains why* the claim about society is true. (Crane 2013: 123)

This quote nicely illustrates the contrast that Crane tries to paint. There are two ways in which one can account for the truth of a sentence without necessarily being committed to what one seemingly talks about. One can either provide a paraphrase that captures the meaning of the original sentences while avoiding any problematic commitments or, alternatively, one can stick with the literal meaning,

and truth, of the original sentence but *explain* how it can be true without being committed to any problematic entity. For instance, a methodological individualist could try to explain why (1) is true while limiting herself to making use only of the actions and intentional states of atomized individuals.

The case of (3) is one that concerns the domain of what exists. The idea is that the literal meaning of true sentences does not always line up with our ontological commitments. When such is the case, we can try to provide a reductive explanation of a truth. According to Crane, this idea straightforwardly also applies to the realm of what does not exist. Literal truths about non-existent objects appear to be problematic because, in virtue of Truth 1, they lead us to Truth 2. But the idea of a reductive explanation of truths can allow us to account for truths about non-existents in a way that does not appeal to Truth 1 and, then, does not lead us to Truth 2. In a nutshell, by making use of the notion of the reductive explanation of a truth, we can try to explain why a truth about a non-existent object is true by appealing only to existent grounds of truth. Accordingly, a sentence like, "John thinks about Pegasus," can be literally true without the need to endorse a Meinongian theory of objects that would supply us with a mind-independent Pegasus as a ground of that truth.

What would a reductive explanation of a sentence like, "John thinks about Pegasus," look like? According to Crane, the existent ground of truth for such a sentence is not hard to find: It is simply the existence of John's intentional experience as of thinking about Pegasus, and that experience's representational content. No more, no less. According to Crane, we must then make a distinction between how the world is and how we represent the world as being. The world is such that there is no relation between John and Pegasus. The only thing there is is an intentional experience of John such that we

describe it as being about Pegasus. But it would be a mistake to discard how we represent the world as being in order to align our theory of intentionality with the way the world is. By first drawing a distinction between substantial and non-substantial properties, and by subsequently reconfiguring the notion of literal truth in order to dispose of Truth 1, we can take at face-value the fact that John is thinking about Pegasus while accounting for this truth merely in terms of John undergoing an intentional experience that is not itself relational. As Crane writes,

> To explain the truths about the non-existent, we need to appeal to facts about the representational content of those truths, or about the ideas they involve. (Crane 2013: 121)

The moral of Crane's account can then be put as follows. A theory of intentionality that wants to make an irreducible use of the notion of intentional object should start with the bedrock claim that some truths about non-existent objects are literally true. It should, however, try to prevent the fact that the endorsement of the possibility of there being such truths creates insuperable philosophical troubles. This can be done, first, by endorsing a distinction between substantial and non-substantial properties and, second, by rejecting an unbridled endorsement of Truth 1.

In fact, there is even more to the theory. According to Crane, we can then use this conception of intentionality to account for more truths about non-existents than truths of intentional sentences. Take, for instance, Pegasus. He does not exist; hence, he cannot possess any substantial properties. There are, however, many truths about him. For instance, he is a mythical character. But that something is a mythical character does not entail that it exists. For something to be a mythical character, it is only required that it be represented in some

way, *i.e.*, in myths and in episodes of thoughts about those myths. The truth of the non-intentional predication, "Pegasus is a mythical character," is then reductively explained in terms of the existence of some representational facts.

In the next section, we discuss how this model of intentionality is supposed to account for the six main features of intentionality that we introduced in Chapter 1 before, in a final section, discussing some objections to Crane's account.

5. Crane's Model and the Six Main Features of Intentionality
5.1 Crane's Model and Aboutness

Crane's model is structurally similar to the IOM, in that it starts from the thought that intentional experiences should irreducibly be described in terms of the notion of intentional object. It is, however, distinct from the IOM because it does not treat intentionality as a relation. Building on the discussion of Chapter 2, Crane's interpretation of the FIS can then be put as follows:

> **Polyadism schema:** "x is about y" amounts to "x stands in a non-substantial relation of 'being about' to an intentional object y, and for some content C, C presents y to x."

As we saw, however, we must distinguish between the claim that this schema gives the *meaning* of the FIS, and the claim that this schema gives us the conditions under which an instance of "x is about y" is true. The only thing that Crane's interpretation gives us is the former. For the latter, we have to resort to the notion of reductive explanation and analyze case by case whether a certain instance of the FIS really obtains.

5.2 Crane's Model and Aspectuality

Crane recognizes the need for the introduction of the notion of content in order to account for aspectuality. His notion of content, however, is not the notion of representational content of the CM, as Crane recognizes that some intentional experiences are non-propositional. Hence, as a general rule, one cannot identify Crane's notion of intentional content with a proposition. According to Crane, generally speaking the notion of content is the way in which an intentional object is presented. Just like in the IOM, Crane's notion of content is not independent of the notion of intentional object. That we need such a notion of content simply follows from the truth that there is more than one way to think about an intentional object.

5.3 Crane's Model and Semantic Normativity

What we said about the IOM will, *mutatis mutandis*, be applicable to Crane's model. That is, some intentional experiences can be semantically evaluated —*e.g.*, John's desire that Federer retires can be satisfied or dissastisfied—because at least some intentional experiences have propositional contents.

5.4 Crane's Model and Intentional Identity

According to Crane's model, just like in the CM and adverbialism but unlike in the IOM, we have to distinguish between two notions of intentional identity: one that applies to cases in which we think about things that exist, and the other to cases in which we think about things that do not exist. When we think about things that exist, we can stand in substantial relation to intentional objects, and such

intentional objects have criteria of identity. On that basis, we can say that two different intentional subjects can be intentionally directed about the same thing.

In a case in which what we think about does not exist, however, the matter is different. Indeed, according to Crane, we do not have the means to say whether my Pegasus is the same as your Pegasus. In fact, we do not have the means to say that the Pegasus I thought about yesterday is the same Pegasus I am thinking about today. Why? Because there is no way to evaluate such judgments without a criterion of identity, and no such criterion of identity is available for non-existent objects.

It is not the time and place to ponder the criterion of identity of existent objects, but it seems quite obvious that for any object that has a criterion of identity, this criterion is not one that depends on representational-dependent properties. On the contrary, we generally say that two different intentional experiences can be said to be about the same object by virtue of the object's substantial properties—its essence, its causal origin, or whatever else counts as relevant to determine its criterion of identity. But since non-existent objects do not have such substantial properties, they cannot have any such criterion of identity.

That said, it certainly seems to us as if we are thinking about the same thing when we are thinking about Pegasus at different times or, say, when we are thinking about a non-existent object that is said to have more than one name, such as, *e.g.*, Hermes and Mercury, who turn out to be the same god. How can we then account for these facts?

According to Crane, we can do so without presupposing numerical identity. Qualitative similarity shall be sufficient. His idea is that when we undergo intentional experiences that are about particular things, either existent or not, we treat information about

them in a particular way. That is, we bundle this information together and do not treat it as determining what our thought is about. In short, we open so-called non-descriptive mental files, as discussed in Chapter 4. We have, for instance, a file for Pegasus that contains information. If our respective files labeled "Pegasus" contain the same information, or a significant overlap of information, we can then be said to think about the same thing. Similarly, if we have two files, one labeled "Hermes" and one labeled "Mercury," and both contain the same information, or a significant overlap of information, then Hermes and Mercury can be said to be the same god.

This appeal to mental files is of course reminiscent of something that we discussed in Chapter 4: the account of thoughts about non-existent particulars offered by a non-descriptivist such as François Recanati, and implemented with the CM. Indeed, according to Recanati, to think about, say, Pegasus is a matter of tokening a non-descriptive mental file that is about Pegasus. As he sees the matter, such thoughts have a singular vehicle, a mental file, but no content.

Crane's appeal to mental files is, however, different from Recanati's. Indeed, in Recanati's account, for a thought to be about, *e.g.*, Pegasus, is for the intentional subject to token a non-descriptive mental file that is about Pegasus. As such, Recanati is faced with the phenomenological problem that his analysis of such thoughts merely seem to amount to what Crane characterizes as "airing an empty vehicle." (Crane 2013: 161) Such thoughts, however, have, just like any other thought, an intentional object and an intentional content. Recanati's account, hence, cannot be correct. As Crane puts it,

> Although I agree with a lot of what Recanati and Taylor say, I cannot put things entirely in their terms. Recanati describes thought-episodes

as 'vehicles' and his view of singular thought-content means that when Le Verrier says to himself 'the discovery of Vulcan will make me famous', the thought has no content. Yet the content of someone's thought is what they are thinking, and how can it be that Le Verrier was not thinking anything, merely airing an empty 'vehicle'? Similarly, Taylor describes a thought's purported singularity as an aspect of the 'form' of a thought rather than its content. But what Le Verrier thinks, I maintain, cannot be characterized simply in terms of its 'form', as if what was going on in him was something without content, in the sense of what is thought. (Crane 2013: 161)

That said, this does not mean that mental files cannot play *another* important role for intentionality. Namely, they play an important, albeit slightly metaphorical, role in explaining how we treat information about particulars, either existent or not. This treatment of information is sub-conscious. That is, whenever we think about, say, Obama we do not bring up all of the information that we possess on Obama. This information, rather, is stored sub-consciously in a file and we can bring it up if such a demand is triggered.

5.5 Crane's Model and Generality

Let us take the following example:

(7) John desires a bottle of wine.

As we saw in Chapter 2, the IOM faces some difficulties in accounting for the truth of such a sentence. Indeed, first of all, there is no indefinite object such as "a bottle of wine" and, secondly, describing John's desire as a desire for an indeterminate object seems

simply mistaken. He wants something from which one can drink and, presumably, one can only drink from a determinate bottle of wine.

The IOM's solution is to interpret (7) by means of a paraphrase. That is, what (7) in fact means is that John desires that he have a bottle of wine. John's desire is then not intentionally directed at a bottle of wine but, rather, at a fact.

According to Crane, however, it is doubtful that this really constitutes a phenomenologically faithful description of John's desire. Indeed, John's desire does not seem to be directed at a fact but, rather, at a bottle, that is, something that is object-like. Indeed, an intentional object, according to Crane, is nothing but something "thrown before experience" (Crane 2013: 92), and, presumably, what is thrown before John's experience is nothing but a bottle of wine. Nothing more, nothing less.

Crane's model, however, can accommodate this fact in a way that is barred to the IOM: Granted that it is true that there cannot be a substantial relation between John and a bottle of wine, there is no reason, however, to deny that there is a non-substantial relation between John and a bottle of wine, granted that this is what is "thrown before" John's experience. As Crane writes,

> My concern… has been to preserve the appearances: to maintain that certain ordinary phenomenological claims are true, while giving an account of them, perhaps in terms that are somewhat different from what these truths might seem to presuppose. In particular, I aim to preserve certain truths about intentional objects. It is literally true that when I want a bottle of inexpensive burgundy, this is the intentional object of my thought. (Crane 2013: 132)

As with non-existence, the key here is to be able to provide a reductive explanation of such literal truths. In the case at hand, this

means being able to explain such a truth in terms of the features of representations themselves. According to Crane, this can be done in a rather standard way, by appealing to the idea that desires have conditions of satisfaction. A desire like John's has conditions of satisfaction with existentially quantified truth-conditions, and this explains why John's desire can be satisfied by any bottle of wine. This is a fact about the nature of John's desire, which exists. But there is a difference between saying that, "John desires a bottle of wine," means the same thing as, "John desires that he have a bottle of wine," and saying that the latter sentence only states the truth-conditions of such a desire, not its content.

5.6 Crane's Model and Non-Existence

We already touched on Crane's solution to non-existence above. That is, a sentence like, "John thinks about Pegasus," can be literally true without there being a relation between John and Pegasus. Such a sentence is instead treated as describing a non-substantial relation between John and Pegasus, and its truth must be explained reductively.

According to Crane, however, we must clearly distinguish between two ways to understand the way John's intentional directedness upon Pegasus could be explained:

(i) "John thinks about Pegasus" is literally true, because Pegasus can be ontologically reduced to the content John's experience.

(ii) "John thinks about Pegasus" is literally true, because the literal truth of this sentence can be explained reductively.

Crane insists that it is (ii) that he is after, not (i). As he sees the matter, it would be mistaken to think that we can ontologically reduce Pegasus to John's experience. Indeed, Crane claims that John's experience could be the same whether or not Pegasus existed. That is, this experience could possess exactly the same intentional properties and would be phenomenally indistinguishable, whether or not Pegasus existed. Both would be about Pegasus, *i.e.*, they would have a content—a Pegasus-content—and an intentional object, namely Pegasus. The non-existent Pegasus and the existent Pegasus, however, cannot be numerically identical.[24] The objection then goes as follows: if there is a claim as to the determination of an intentional object by the content of its intentional experience, then two token of a same kind of intentional experience should determine the same intentional object. But the non-existent and the existent Pegasus cannot be the same object. Hence, (i) cannot be true. As Crane puts it:

> To explain the relationship between truths about non-existing intentional objects and existing things, the notion we need is not supervenience, but *explanatory reduction*: we need an explanation of why it is true that [John thinks about Pegasus], which appeals only to what does exist. So I will not talk of 'ontological reduction'… since I reserve this phrase for the identification of a domain of entities as a sub-domain of another. And non-existent objects are not, of course, entities of any kind. (Crane 2013: 134)

What then would explains the truth of "John thinks about Pegasus?" Well, according to Crane, the answer to that question is pretty simple. "Thinking about" is a non-substantial relation. Hence, John can think

---

[24] Note that they could be if we were able to endorse something like a Meinongian doctrine of natures, but we rejected such a theory in Chapter 3.

about Pegasus even if Pegasus does not exist. The only thing that is required for "John thinks about Pegasus" to be literally true is an intentional experience as of being directed upon Pegasus. As a result, the non-substantial relation of "thinking about" obtains between John and Pegasus, and Pegasus instantiates the non-substantial property of being an object of thought.

In that sense, since, "John thinks about Pegasus" is literally true, we are allowed to infer, unlike in the CM and in adverbialism, that there is something about which John is thinking, namely Pegasus. In that sense, Crane's model cannot be accused of making use of a notion of intentional directedness that is subject to success or failure.

That said, one might wonder what are the conditions under which an existentially quantified sentence such as, say, "There is something about which John is thinking" is true. Indeed, in the case of the IOM, the truth-conditions of such a sentence were clearly statable. Namely, the Meinongian theory of objects fills up the overall domain of quantification with so-called pure objects, in a way that is independent of us. A sentence such as, "There is something about which John is thinking" can then said to be true *iff* (i) John is indeed thinking about something and (ii) there is an element of the overall domain such that it can serve as a witness of this quantified sentence. The Meinongian theory of objects, however, guarantees that if condition (i) is met, then condition (ii) is also met. Crane, however, rejects the idea of a Meinongian theory of objects. What, therefore, does account for the truth of a sentence like, "There is something about which John is thinking?"

What we said above already provided us with a hint of Crane's answer. By virtue of the fact that a non-substantial relation obtains when John is thinking about Pegasus, Pegasus instantiates the property of being an object of thought. According to Crane, then, a true quantified sentence such as "There is something that John is

thinking about" can be validly inferred and from a sentence like "John thinks about Pegasus" because there is an object of thought, Pegasus, that can serve as a witness for the quantified expression: "∑(x) Think about (John, x)."

According to Crane, however, such a solution enjoins us to reject a classical conception of domains of quantification. He summarizes this conception as follows:

> This is one way to describe my departure from the standard view: for many who hold the standard view want to explain representation in terms of an antecedent conception of domains of quantification, and relations defined on these domains. (Crane 2013: 41-42)

The idea is the following. A traditional understanding of domains of quantification takes them to be sets of *entities*. These entities may all exist, as in classical logic, or may not exist, as in a Meinongian outer-domain of quantification. The key idea, however, is that members of domains are entities and that we understand intentional experiences in terms of these entities—as well as relations to them—, and not the other way around, *i.e.*, we do not understand domains—as well as some predicative operations over them—in terms of intentional experiences.

According to Crane, we should reject this traditional conception and replace it by the idea that what is relevant for the evaluation of quantified sentences such as: "∑(x) Think about (John, x)" are not entities but, rather, objects of thoughts, that is, these things that are relevant for the description of the way we think and talk. As he puts it:

> In one traditional terminology, the domain of quantification was called the universe of discourse. This term gives a hint as to how we should

think of quantification metaphysically if we are going to make literal sense of sentences [that existentially quantify over things that do not exist]. The universe of discourse contains all the items we assume or stipulate to be relevant to our discourse. An item here is simply something which can be thought or spoken about: an object of thought or discourse, in the sense I introduced above. The domain of quantification consists of just those objects of thought relevant to the truth or falsehood of the quantified claim. (Crane 2013: 38)

As such, then, if John is thinking about Pegasus, then we can infer that there is something that John is thinking about, namely Pegasus. Hence, we can say that Pegasus can serve as the value of the variable in: "$\sum(x)$ Think about (John, x)."

However, as we will see in the upcoming critical part of this chapter, such a conception of domains of quantification is not without problems.

6. Some Questions About Polyadism

6.1 Circularity

An adequate theory of intentionality should be one that is able to account for the feature that we called "non-existence." As we just saw, Crane thinks that he has provided an account that meets this demand. For one to think about a non-existent intentional object is for one to stand in the non-substantial relation of "being about" to a non-existent intentional object. In other words, if Crane is correct, then a sentence describing an intentional experience such as

(8) John thinks about Pegasus

can be literally true because there is a non-substantial relation between John and Pegasus.

But what are the conditions under which (8) is true? According to Crane, we do not necessarily need to look for a ground of truth that shares exactly the shape of the proposition expressed by (8). What we should be looking for is merely the fact that John is undergoing an intentional experience that we may describe as making (8) true.

How, however, should one spell out the conditions under which (8) is true? Well, a good candidate seems to be the following:

(9) John thinks about Pegasus.

But this look quite circular, does it not? Saying that (9) explains (8) does really look like (8) is explaining itself. How can this constitute a satisfying account of non-existence?

Crane explicitly recognizes that his account is circular, but takes it as meaning one thing: The notion of representation must be

taken as basic. In other words, there is no way to understand how John can think about Pegasus can be true but in the terms of the truth of, "John thinks about Pegasus." That is, we cannot explain this fact about John's representational powers in terms of a causal relation between John and Pegasus or a non-causal relation between John and a proposition that represents Pegasus. If John thinks about Pegasus, then "John thinks about Pegasus" is literally true because of the fact that John is undergoing an experience to be described in just that way is a basic truth about representation.

That said, it is important that we do not appeal to non-representational facts to explain this feature of representations and, in particular, that we do not appeal to there being mind-independent non-existent intentional objects that our experiences can somehow target. On the contrary, there is a non-existent Pegasus because it has the non-substantial property of being thought about by John, a property that it can have indifferently to its ontological status. As Crane puts it,

> This is not really a vicious circle, since what I mean by 'understanding the non-existent' really amounts to understanding truths, understanding how they can be true, and understanding what makes them true. What makes these truths true are certain facts about representation. (Crane 2013: 168)

As Crane sees the matter, then, his account is only available to those who are ready to regard representation as being basic. This is a price that one needs to be ready to pay to be able to account for the phenomenological fact that (8) can be literally true without endorsing some form of Meinongianism.

6.2 Phenomenological Adequacy

In their respective reviews of Crane (2013), Anthony Everett and Michelle Montague independently point—quite ironically—to a concern with respect to the phenomenological adequacy of Crane's model. Here is Montague's version of the objection:

> Tom's thought about Vulcan is made true by Tom's thought about Vulcan. Yes, but what does that truth-maker consist in? Asking this may provide another way of trying to understand, perhaps indirectly, what Tom's thought about Vulcan consists in. Vulcan, in so far as it has properties, has properties like *being a postulated planet*. So is Tom's thought about a postulated planet? Surely that is not how it seems to Tom (suppose he knows it's postulated and thinks the postulation is also true). For Tom, his thought seems to be about a planet. Note, our thoughts often get things wrong about the world, but in the case of Tom, according to Crane's theory, he's wildly wrong about the object of his thought. Imagine someone asking me what I'm thinking about and I say, "I'm thinking about a mountain, you know, the golden mountain" and I'm met with the response "You're definitely not thinking about a mountain". The question is, how seriously do we take our claims about what we are thinking about? (Montague 2015: 337)

The objection seems to be the following. In a case like the one of Tom's thought about Vulcan, Tom is thinking about something, Vulcan, that possesses the property of being a postulated planet. Hence, it is true that Tom is thinking about a postulated planet. Tom, however, would not describe his intentional experience as being about a postulated planet but, rather, as being about planet. Hence, the objection goes, how are we to reconcile these two claims? Should

not we say that what John is in fact thinking about is a planet, and not a postulated planet?

And here is Everett's version:

> My second reservation concerns whether Crane can provide an adequate account of all talk and thought which purports to be about the non-existent. For Crane, while non-existent objects may have pleonastic properties, they cannot have genuine substantial properties. Vulcan has various pleonastic properties, such as being thought about by Le Verrier, and being a postulated planet. But Vulcan cannot have the property of being a planet. The property of being a planet is a substantial property and only existing things can have substantial properties.

> However suppose that Le Verrier has just posited Vulcan. Consider the following report:

> (10) Le Verrier is thinking about a planet.

> Intuitively (10) seems correct. However, for Crane, while it is correct to say Le Verrier's thought has an object, Vulcan, this object is not a planet. So on Crane's account (10) should count as false, for (10) characterizes Le Verrier as thinking about a planet when he is not.

> These sorts of cases bring out the fact that we may report thoughts 'about' the non-existent from two perspectives. We may report them from an external perspective in which we explicitly mark the things they are 'about' as mythical or fictional or otherwise non-existent, as we do with (11):

> (11) Le Verrier is thinking about a postulated planet.

> However it is usually far more natural to adopt an internal perspective where we talk and think as if thoughts 'about' the non-existent were directed towards real objects which have substantial existence entailing properties. I suggest this is what is going on with (10). We talk and think as if, we make-believe that, the world contains not merely the things that really exist but also a planet towards which Le Verrier's thought is directed. (Everett 2015: 1276)

Everett's version, contrary to Montague, explicitly brings in a distinction between two perspectives: an internal and an external one. The idea is that from an internal perspective, Tom's—or Le Verrier's —thought about Vulcan is about planet. Hence, from the internal perspective, Everett's (10) is true. However, from an external perspective, *i.e.*, from the perspective of someone who knows that Pegasus does not exist, then (10) is false and must be replaced by Everett's (11). Now, Everett's point is slightly different from Montague's. Namely, he claims that the problem with Crane's account is that it is more plausible to describe his (10) as being a case in which John engages into a make-believe activity that concerns a *real* planet rather than as being a case in which John thinks about a non-existent postulated planet.

The two objections are very close but most nonetheless be answered separately. We suggest that Montague's point can be answered merely by appealing to the distinction between intentional object and intentional content. Indeed, a same intentional object can be presented under different contents and there is no reason to be alarmed by the fact that a content presents an object in a wildly mistaken manner. To answer Montague's question, one could simply answer: "Look, I know that the golden mountain is a postulated planet. However, I can think about the golden mountain *as* a planet

and *as* a postulated planet. When I am telling you that I am thinking about the golden mountain, I am not fooling myself into thinking that I am thinking about a real golden mountain. I know that there is no such mountain, and that what I am intentionally directed upon merely is a postulated planet. But right now I am thinking about it as a planet. Why? Maybe because of the kind of imaginative project I am engaged in. That is, I am imagining that the golden mountain is a planet because, say, I am imagining climbing it. And one can only climb a real planet, not a postulated one."

At this point, however, it seems that Everett's point kicks in. Indeed, one might well take in this answer to Montague's objection but answer, alongside Everett, that in such a case one would be better off with describing one's intentional experience in terms of a make-believe activity of thinking about a real planet rather than as an intentional experience directed upon a postulated planet. What should a proponent of polyadism say at this point?

Three elements of answer can be mentioned. The first one is that it does not seem that Everett's objection makes polyadism extensionally inadequate. In other words, it does not seem that Everett's objection makes it such that there are instances of intentional experiences that polyadism cannot account for. Indeed, the distinction between intentional content and intentional object is sufficient to account for the distinction between (10) and (11). Everett, however, concedes this point as his objection only appeals to idea that it is "far more natural" to describe the difference between his (10) and (11) in his terms rather than in polyadism's ones. His claim, hence, is really one of naturalness and not of extensionality. Everett certainly agrees with that, but this shall turn out important for what comes next.

The second point is that, contrary to polyadism, Everett's account does seem to be extensionally inadequate, and on two

different counts. First, according to Everett, all truths about non-existents—including truths like his (10), but also non-intentional ones—must be explained in terms of make-believe. For instance, if "Sherlock Holmes lives on 221B Baker Street" is true, then it is true only because we make-believe that it is true, *i.e.*, it is true under a certain intensional embedding, such as: "According to the Holmes fiction, Sherlock Holmes lives on 221B Baker Street." Such a truth, however, is not a literal truth. It is, rather, only a "make-believe truth," *i.e.*, a truth that is true in virtue of some make-believe activities.

Take, however, a sentence such as "Sherlock Holmes is more famous than any living detective; for example, Sherlock Holmes is more famous than Sir Ian Blair" (Crane 2013: 135). This sentence is true. However, can we find an intensional embedding that makes it true? That is, is there an operator such as: "According to the fiction, …" that make it true? It does not seem to be the case as there is no fiction or anything similar of which both Sherlock Holmes and Sir Ian Blair are characters that could make it true. Hence, Everett's strategy cannot be extensionally adequate as a general way to account for the truth about non-existents.

Of course, one may try something along the following lines: "According to what people say, Sherlock Holmes is more famous than Sir Ian Blair" or "According to what people think, Sherlock Holmes is more famous than Sir Ian Blair." This, however, would get things the wrong way around. Indeed, that Sherlock Holmes is more famous than Ian Blair is not a truth about what people say or people think. It is not like, say, the truth that people generally say that cold water cools down hot water, even if we know that the contrary is true. Indeed, that Sherlock Holmes is more famous than Sir Ian Blair is a genuine truth about the world, not about what people say or think. That said, it is nonetheless true that it is the case that Sherlock

Holmes is more famous than Ian Blair because of the way people talk and think. Indeed, it is because people talk and think more about Sherlock Holmes than Ian Blair that the former is more famous than the latter. But this is a fact that speaks in favor of polyadism, not of Everett's account.

Second, Everett must claim that if true, "John thinks about Pegasus" can be true only by virtue of some kind of make-believe activity. This, however, seems to get something wrong about the phenomenology of such intentional experiences. The make-believe strategy should analyze this sentence along the following way: "According to John's make-believe activity, John is thinking about Pegasus." But this would claim that, when John is thus thinking about Pegasus, he is not, as a general rule, really trying to engage with our world but, rather, with the world of his make-believe activities. This seems wrong. At least in some cases, thinking about Pegasus should indeed count as a way to try to engage with our world, and not with something else.

Accordingly, we can return the charge of phenomenological adequacy back to Montague and Everett: it might sound weird that intentional subjects, purportedly or not, sometimes engage with intentional objects whose nature they deeply misconceive, as in Montague's above objection. But it seems even weirder to claim that such cases should be analyzed as engagement with objects that really have the properties represented in the experience, though only according to some make-believe activity, as in Everett's alternative proposal.

6.3 Existent Intentional Objects

Presumably, at least some of the things about which we think exist. How does Crane's model account for these cases? Are such cases also

instances of non-substantial relations? If so, how do they manage to hit their existent targets?

Crane's answer is quite simple: Some intentional experiences are instances of substantial relations. For instance, one may think about Obama and manage not only to think about Obama, but also to *refer* to him, where reference is to be understood as a real relation. But one should distinguish between two questions:

(a) If, *e.g.*, John manages to refer to Obama, does reference serve as a constitutive condition on John's experience being intentional?

(b) If, *e.g.*, John manages to refer to Obama, does reference help determine what is the object of John's intentional experience?

According to Crane, it is essential to distinguish between these two questions. Indeed, one can say that the answer to (b) is positive while the answer to (a) is negative. In other words, that John stands in a substantial relation to Obama does not play any constitutive role on it being an intentional experience, defined as an experience having an intentional object. Were this relation not to obtain, John's experience could still be said to have an intentional object.

What is important, according to Crane, is that we realize that there is more than one way for an intentional experience to have an intentional object. One of them is to be a substantial relation. But this is only *one* way for an intentional experience to have an intentional object. Another way is simply to be a non-substantial relation. We should simply not hope for a systematic answer to the question, "How does an intentional experience get its intentional object?" He explains this as follows:

There are many varieties of intentionality, and we should not expect there to be one account that unifies all these different varieties, which will explain, for any kind of mental representation whatsoever, why it is that some object rather than another is the real object of a given intentional state. Intentionality is irreducible. (Crane 2013: 115)

That said, one might be willing to ask a further question. In a counterfactual situation in which Obama does not exist, can we say that John's intentional experience of thinking about Obama in our world and the one in the counterfactual situation are about the very same object? As we saw above, Crane denies that this answer can be positive. Indeed, non-existent intentional objects cannot flank non-trivial identity sentences. Take the following identity sentence, where "Obama" stands for Obama and "Obama$_w$" stands for John's intentional object in the counterfactual situation: "Obama is a human being, and has parents. Obama$_w$, on the other hand, is not a human being, and does not have parents." These two things cannot, then, be identical. The two thoughts, however, have an intentional object.

6.4 Meaning

Another concern focuses on Crane's account of the meaning of the proposition expressed by a sentence like (8). According to him, it can be modeled as "Think about (John, Pegasus)," where "thinking about(x, y) should merely be conceived as a non-substantial relation. Moreover, he is pretty explicit about the conditions that a theory should meet in order to be able to account for the semantic content of such a sentence.

First, since this theory should be able to account for meaningful non-referential singular terms, it should not be cashed out

in terms of a classical logic. The logic that we need here is a free logic. Second, since such a proposition is taken to be true, then the free logic that we need cannot be a negative one. Indeed, according to negative free logic, all simple sentences formed by means of a non-referential singular term turn out to be meaningful but false. Hence, negative free logic cannot account for the fact that "Think about (John, Pegasus)," is true. The logic that we need, then, should be a positive free logic, *i.e.*, one that accounts for true simple sentences involving non-referential singular terms.

Beyond the mere logical requirements, it is also quite clear that the semantic content of such a proposition can be neither an existentially quantified content nor a Millian content. Indeed, if the former were the case, then (8) could not have the literal meaning Crane claims it possesses and, similarly, if the latter were the case, then presumably there should be Meinongian non-existent objects or, worse, we should give up on a great scale to the feature we called non-existence.

In sum, then, Crane is committed to a non-descriptive, non-Millian semantic account of a proposition like the one expressed by (8), to be couched in a positive free logic, but he does not tell us much more than that. His account dramatically lacks any positive claim with respect to semantic content. That we can really understand the content of such propositions is then conditional on the availability of a systematic account for the semantic content of such propositions.

6.5 Quantification

Crane directs the following objection against his understanding of domains of quantification:

It might be objected, however, that thinking of the members of a domain of quantification as objects of thought gives rise to paradox. An object of thought is just something thought about. But surely we can quantify over things that have never been thought about: for example, we can say 'some things have never been thought about'. This surely must be true. But if so, how can the domain of quantification consist of objects of thought? (Crane 2013 : 39)

The objection can be reconstructed as follows:

1) The members of the domain of quantification are objects of thought;
2) The sentence : "Some things have never been thought about" is true;
3) If this sentence if true, then there are objects that can make an existentially quantified   sentence true and yet they are not objects of thoughts;
4) Hence, *contra* 1), there are members of domain of quantification that are both objects of thought and not objects of thought.

According to Crane, however, this objection does not go through because 3) is false. Indeed, that some objects have never been thought about does not mean that they are not objects of thought. It merely mean that they are objects of thought that have never been thought about *except* by means of an existentially quantified thought with the content:

(10) Some things have never been thought about.

As he sees the matter, what needs to be endorsed in order to avoid this objection is the claim that quantifying over something is a way of thinking about something. As he puts it :

> I can use a quantified noun phrase to 'talk about' things, even if those things cannot be talked about in other ways. It might seem paradoxical to say 'some things have never been talked about', but if we agree that quantified noun phrases are ways of talking about things, then we should understand this as conveying the following: some things have never been talked about except by being talked about in this way. This is comparable to what one should say to Berkeley when he says that one cannot conceive of an unconceived tree. Of course, by conceiving of a tree as unconceived, what one means is that it is not conceived in any way other than in this act. If there is a paradox here, it is not one which is specific to the view of quantification defended here. (Crane 2013 : 39)

Let us assume that Crane's rejoinder to this objection is correct. Does it mean that his account of domains of quantification is without further problem? One further concern that may be raised goes as follows. Take the following sentence:

> (11) Every things have been thought about[25].

It should be pretty obvious that (11) is false. Indeed, there are many things that have not been thought about—recalling, *e.g*, Pascal's double infinity, it is in fact highly probable that there are more things that have not been thought about than things that have been thought

---

[25] Note that Crane explicitly endorses a binary reading of the quantifiers in Chapter 2 of Crane 2013.

about. According to Crane's account of domains, however, (11) turns out to be trivially true. Indeed, if quantifying is way of thinking about something, then quantifying over everything is a way of thinking about everything. Hence, by thinking a thought with (11) as content, then one manages to make it true. But (11) is clearly false. Hence, Crane's account must be false.

Of course, Crane possesses an easy way out of this objection. Namely, he could simply claim that (11) is ambiguous, and true according to one reading and false according to another one. A first difficulty is that it does not seem that (11) is ambiguous. Let us, however, admit that this is the case. According to this rejoinder, (11) would be presumably be true when interpreted as ranging over objects of thoughts. Indeed, trivially, all objects of thoughts in Crane's sense can be said to have been thought about. Hence, under that first reading, (11) is true. But what about the reading that turns (11) into a falsity? It seems that the content of this second reading should rely on the truth of: "Some things have not been thought about." But since, as he puts it, "Some things have not been thought about" is way of thinking about things, then it would follow that, in fact, every things have been thought about. Hence, both readings would in fact make (11) true.

7. Conclusion

In this last chapter, we discussed what we called polyadism. Recently defended by Tim Crane, this model tries to account for the phenomenological fact that each intentional experience must be described in terms of the notion of intentional object without thereby endorsing some form of Meinongianism.

As we saw, this requires that three main theses be endorsed:

(T13) A theory of quantification is not committed to an objectual reading of quantifiers;

(T14) For a restricted range of non-existence entailing properties, Meinong's Principle of Independence of *Sosein* from *Sein* holds;

(T15) Some sentences of the form "*a* is *F*" can be literally true, without there being any ground of truth for such sentences to the form, "*a* that is *F*."

(T13) was endorsed on the basis of the discussion of quantification provided in Chapter 3. (T14) was defended in terms of a distinction between two kinds of properties: substantial and non-substantial properties. The only truths that hold of non-existent objects, with the exception of negative existentials and trivial identities, are truths that result from the predication of a particular kind of non-substantial properties: representation-dependent one.

(T15), on the other hand, results from the rejection of a principle we called Truth 1.

> **Truth 1:** If a sentence of the form "*a* is *F*" is literally true, then the ground for this truth must be that there is some object, *a*, that instantiates *F* (and something similar holds for *polyadic* predications).

This principle was rejected on the ground that we must distinguish between two ways in which a sentence can be literally true: either it is true by virtue of the conditions stated in Truth 1, or it is true by virtue of a reductive explanation, where this is to be distinguished from a paraphrase.

The rejection of Truth 1, then, paved the way for the rejection of Truth 2.

> **Truth 2:** If *a* does not exist, then this ground for truth mentioned in Truth 1 must be partly constituted by a non-existent object, *i.e.*, *a.*

As such, then, there can be truths about the non-existent, intentional, and non-intentional truths, all to be explained in terms of existent representational facts.

Crane's proposition, we claimed, allows us to substantially reconfigure the logical space of the theories of intentionality. Indeed, with the availability of the polyadism's schema as an interpretation of the FIS, we are able to distinguish between two different issues. First, whether a certain theory of intentionality is representational. Second, whether a certain theory of intentionality uses a notion of aboutness that is not subject to success or failure. Without an interpretation of the FIS like polyadism's schema, it seems that the price to pay for defending a theory of intentionality that uses a notion of aboutness that is not subject to success or failure is to endorse a relational theory of intentionality. The endorsement of polyadism schema, however, allows one to defend a theory of intentionality that uses a notion of aboutness that is not subject to success or failure without thereby collapsing into some form of Meinongianism. This is a new and interesting development for the theory of intentionality.

However, we saw that polyadism is not without problems of its own. First, it should be complemented by an explicit account of the semantic content of propositions expressed by literal truths such as: "John thinks about Pegasus." Another worry concerned Crane's interpretation of domains of quantification, as we saw that his account made a sentence like: "Every things have been thought about" trivially true.

Concluding Remarks

 In this dissertation, our general aim was to inquire into the way one should describe what sentences such as "John thinks about Pegasus" or "John thinks that Pegasus is a horse" describe, namely—as we defended the claim in Chapter 1—intentional experiences.

Our starting hypothesis was that all intentional experiences are instances of the fundamental intentional schema (FIS): "*x* is about *y*." We claimed that the task of a theory of intentionality should be to provide an interpretation of this schema in such a way that it can be taken to account for six main features of intentionality: aboutness, aspectuality, semantic normativity, intentional identity, generality, and non-existence.

The first interpretation of that schema we focussed on, discussed in Chapter 2, was the one of what we called the intentional object model (IOM), namely:

> **IOM schema:** "*x* is about *y*" amounts to "*x* stands in the relation of 'being about' an intentional object *y,* and for some content *C, C* presents *y* to *x*."

As we saw, such an interpretation of the FIS must be wedded to a Meinongian theory of objects since it must account for the following two theses:

(i) Some intentional objects do not exist; and:

(ii) Non-existent intentional objects can be the relata of the relation of "being about."

A useful slogan to characterize this first interpretation could be the following: "In all cases, explain how an intentional experience can be intentionally directed upon something in terms of truths about mind-independent intentional objects." Indeed, according to the IOM, the key to an adequate theory of intentionality is that it should be conceived within the (rather generous) boundaries of a Meinongian theory of objects according to which there are, in a completely mind-independent way, more things than those that exist. However, we argued in Chapter 3 that there are good grounds not to commit oneself to Meinongianism and, consequently, that this first interpretation of the FIS does not seem to be an adequate one.

In Chapter 4, we then moved to a second interpretation of the FIS that aimed to avoid the problems faced by the IOM schema by providing a different kind of interpretation of the FIS, namely a paraphrastic one:

> **The CM schema**: "x is about y" amounts to "x represents y"; and x can be said to represent y just in case there is a representational content C that either is, or could be, true of y.

A paraphrase denotes a relation between two things, an original sentence S and a paraphrase S' of S. As we saw, there are two ways to conceive of the relation between S and S'. One can either try to use a paraphrase S' to clarify the natural meaning of S or, alternatively, one can try to use a paraphrase S' to align the meaning of S with independently determined theoretical commitments. We interpreted the CM as trying to provide an interpretation of the first kind. The resulting theory can be summarized by means of the following slogan: "In all cases, explain how an intentional experience can be intentionally directed upon something in terms of the semantic properties of its representational content."

We then argued that neither does the CM schema constitute a good interpretation of the FIS. First, we argued that a theory of intentional experiences should countenance the existence of both objectual and propositional intentional experiences. The CM schema, however, applies only to intentional experiences with a propositional representational content. Hence, it is extensionally inadequate. Second, its analysis of aboutness is not strong enough to account for the intuitive notion of aboutness that characterizes the FIS.

In Chapter 5, we turned to a third option, adverbialism, as recently defended anew by Uriah Kriegel. Unlike the IOM, but like the CM, adverbialism offers a paraphrastic analysis of the FIS along the following lines:

Adverbialist schema: "*x* is about *y*" amounts to "*x* is intentionally directed *y*-wise."

A slogan we could use to summarize this schema could be the following: "In all cases, explain how an intentional experience can be intentionally directed upon something in terms of adverbial modifications of consciousness."

This third schema presents some advantages over the previous ones. Indeed, despite the fact that adverbialism eschews talking in terms of intentional experiences being intentionally directed upon intentional objects, Kriegel's use of the idea of foreignness constitutes a step in the direction of uniformly describing all intentional experiences in terms of "the presence to the mind—or apparent presence to the mind—of things in the world" (Crane 2013: 4).

There are nonetheless good reasons to believe that the adverbialist schema does not constitute an adequate interpretation of the FIS. First, the determinable-determinate framework that is used by Kriegel to avoid standard objections to adverbialism saddles intentional experiences with

too much contents than required. Second, the notion of phenomenal foreignness is not strong enough to properly delineate the extension of the class of all intentional experiences, as some experiences instantiate foreigness, though they only doubtfully could be dubbed "intentional."

Finally, we turned to a last interpretation of the FIS, the one provided by polyadism. According to polyadism, the IOM's appeal the notion of intentional object as a necessary conceptual component of an adequate description of intentional experiences is fundamentally correct. However, what is required for a theory of intentionality that irreducibly appeals to such a notion is to be able to disentangle this claim from the IOM's endorsement of Meinongianism. As we saw, Crane proposes to do just this in terms of the following schema:

> **Polyadism Schema:** "*x* is about *y*" amounts to "*x* stands in a non-substantial relation of 'being about' to an intentional object *y*, and for some content *C*, *C* presents *y* to *x*."

According to Crane, the key idea to understand how this schema can be said to universally apply to intentional experiences without presupposing Meinongianism is that we shall primarily focus on the notion that all instances of the FIS—including those that describe intentional experiences, either propositional or not, that are about things that do not exist—are literally true. Second, the notion of literal truth must be divorced from the idea that if a sentence of form "*a* is *F*" is literally true, then there must be a ground of truth for that sentence that instantiates what Crane metaphorically calls a "similar shape."As such, then, polyadism can be summarized by means of the following slogan: "In all cases, explain how an intentional an intentional experiences can be intentionally directed upon something in terms of the literal truths of sentences used to describe intentional experiences—and, crucially,

distinguish between a literal truth and the shape of the reality described by that truth."

As we argued, Crane's proposal allows us to reconfigure the logical space of theories of intentionality in a substantial way. Indeed, his account allows to distinguish between, on the one hand, relational and representational theories of intentionality and, on the other hand, theories of intentionality that use a notion of aboutness submitted to conditions of success and failure and theories that do not. His development of polyadism allows a representational theory of intentionality to be one that, contrary to the CM and adverbialism, does not make use of a notion of aboutness that is subject to conditions of success and failure.

Crane's account, then, seems to offer us quite a lot. As such, it seems to go a long way towards allowing us to make sense of the following passage of Husserl's fifth *Logical Investigations*:

> If I represent the god Jupiter, this god is my presented object, he is 'immanently present' in my act, he has 'mental inexistence' in the latter, or whatever expression we may use to disguise our true meaning. I represent the god Jupiter: this means that I have a certain presentative experience, the presentation-of-the-god-Jupiter is realized in my consciousness. This intentional experience may be dismembered as one chooses in descriptive analysis, but the god Jupiter naturally will not be found in it. The 'immanent', 'mental object' is not therefore part of the descriptive or real make-up (*deskriptiven (reellen) Bestand*) of the experience, it is in truth not really immanent or mental. But it also does not exist extramentally, it does not exist at all. This does not prevent our-idea-of-the-god-Jupiter from being actual, a particular sort of experience or particular mode of mindedness (*Zumutesein*), such that he who experiences it may rightly say that the mythical king of the gods is present to him, concerning whom there are such and such stories. If, however, the intended object exists,

nothing becomes phenomenologically different. It makes no essential difference to an object presented and given to consciousness whether it exists, or is fictitious, or is perhaps completely absurd. I think of Jupiter as I think of Bismarck, of the tower of Babel as I think of Cologne Cathedral, of a regular thousand-sided polygon as of a regular thousand-faced solid. (Husserl 2002: 99—modified translation, see footnote 23 below).

If one regards this passage of Husserl as providing us a particularly sharp description of the general nature of intentional experiences, then Crane's interpretation of the FIS turns out to be a particularly interesting one. If one does not regard this passage of Husserl in this way, then maybe a second look would help.

However, we saw that Crane's account is not without problems. We mentioned a handful of issues, and several others certainly loom around. Further work, hence, is require to complete and develop the view. Asked to answer the question: "What's philosophy?", Anthony Appiah answers: "Everything is much more complicated than you first thought." After these six chapters, we cannot say that we just started to think about intentionality. What seems pretty obvious, however, is that

intentionality is in fact much more complex that what we thought so far.[26]

---

26 I modify Findlay's translation of the *Logical Investigations* which I find faulty. In the original, Husserl writes: "Stelle ich den Gott Jupiter vor, so ist dieser Gott vorgestellter Gegenstand, er ist in meinem Akte "immanent gegenwärtig", hat in ihm "mentale Inexistenz", und wie die in eigentlicher Interpretation verkehrten Redeweisen sonst lauten mögen. Ich stelle den Gott Jupiter vor, das heißt, ich habe ein gewisses Vorstellungserlebnis, in rmeinem Bewußtsein vollzieht sich das Den-Gott-Jupiter-Vorstellen. Man mag dieses intentionale Erlebnis in deskriptiver Analyse zergliedern, wie man will, so etwas wie der Gott Jupiter kann man darin natürlich nicht finden; der "immanente", "mentale" Gegenstand gehört also nicht zum deskriptiven (reellen) Bestande des Erlebnisses, er ist also in Wahrlheit gar nicht immanent oder mental. Er ist freilich auch nicht extra mentem, er ist überhaupt nicht. Aber das hindert nicht, daß jenes Den-Gott-Jupiter-Vorstellen wirklich ist, ein so geartetes Erlebnis, eine so bestimmte Weise des Zumute- seins, daß, wer es in sich erfährt, mit Recht sagen kann, er stelle sich jenen mythischen Götterkönig vor, von dem dies und jenes gefabelt werde. Existiert andererseits der intendierte Gegenstand, so braucht in rphänomenologischer Hinsicht nichts geändert zu sein. Für das Bewußtsein ist das Gegebene ein wesentlich Gleiches, ob der vorgestellte Gegenstand existiert oder ob er fingiert und vielleicht gar widersinnig ist. Jupiter stelle ich nicht anders vor als Bismarck, den Babylonischen Turm nicht anders als den Kölner Dom, ein regelmäßiges Tausendeck nicht anders als einen regelmäßigen Tausendjlächner." (Husserl, *Logische Untersuchungen*, V, §11)
        Findlay, however, translates as: "If I have an idea of the god Jupiter, this god is my presented object, he is 'immanently present' in my act, he has 'mental inexistence' in the latter, or whatever expression we may use to disguise our true meaning. I have an idea of the god Jupiter: this means that I have a certain presentative experience, the presentation-of-the-god-Jupiter is realized in my consciousness. This intentional experience may be dismembered as one chooses in de- scriptive analysis, but the god Jupiter naturally will not be found in it. The 'immanent', 'mental object' is not therefore part of the descriptive or real make-up (deskriptiven reel/en Bestand) of the experience, it is in truth not really immanent or mental. But it also does not exist extramentally, it does not exist at all. This does not prevent our-idea-of-the-god-Jupiter from being actual, a particular sort of experience or particular mode of mindedness (Zumutesein),

Bibliography

Azzouni, Jody (2010). *Talking about Nothing: Numbers, Hallucinations, and Fictions*. Oxford: Oxford University Press.

Azzouni, Jody (2012). "Referring to What Is and to What Isn't." In: Kabasenche, O'Rourke, and Slater (2012).

Bayne, Tim, and Michelle Montague (2010). *Cognitive Phenomenology*. Oxford: Oxford University Press.

Beckermann, Ansgar & McLaughlin, Brian P. (eds.) (2007). *Oxford Handbook to the Philosophy of Mind*. Oxford: Oxford University Press.

Bell, David (1990). "How 'Russellian' Was Frege?" In: *Mind* 99 (394): 267-277.

Bennett, Jonathan (1974). *Kant's Dialectic*. Cambridge: Cambridge University Press.

Berto, Francesco (2013). *Existence as a Real Property: The Ontology of Meinongianism*. Dordrecht: Springer.

Berto, Francesco, and Matteo Plebani (2015). *Ontology and Metaontology: A Contemporary Guide*. London: Bloomsbury Academic.

Block, Ned (1995). "On a confusion about a function of consciousness." In: *Brain and Behavioral Sciences* 18 (2): 227-–247.

Bradon-Mitchell, David & Jackson, Frank (1996). *Philosophy of Mind and Cognition*. Oxford: Blackwell.

Brentano, Franz (2015). *Psychology from an Empirical Standpoint*. London: Routledge.

Bricker, Phillip (2016). "Ontological Commitment." In: *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), forthcoming URL = <https://plato.stanford.edu/archives/win2016/ entries/ontological-commitment/>.

Burge, Tyler (1974). "Truth and singular terms." In: *Noûs* 8 (4):309-325.

Byrne, Alex (2006). "Intentionality." In: Sakar, Sahotra & Pfeifer, Jessica (eds.)

Carnap, Rudolf (1956). *Meaning and Necessity.* Chicago: University of Chicago Press.

Carruthers, Peter & Veillet, Bénédicte (2011). "The Case Against Cognitive Phenomenology." In: Bayne, Tim and Montague, Michelle (ed.) (2011).

Cartwright, Richard L. (1960). "Negative existentials." In: *Journal of Philosophy* 57 (20/21):629-639.

Chalmers, David (ed.). (2002). *Philosophy of Mind. Classical and Contemporary Readings.* Oxford: Oxford University Press.

Chalmers, David, Manley, David & Wasserman, Richard (eds.) (2009). *Metametaphysics: New Essays on the Foundations of Ontology.* Oxford: Oxford University Press.

Chisholm, Roderick M. (1960). *Realism and the Background of Phenomenology.* Glencoe, Ill.,Free Press.

Chrudzimski, A. (2014). "Intentional Objects and Mental Contents." In: Fréchette, G. (ed.) 2014.

Chudnoff, Elijah (2015). *Cognitive Phenomenology.* London: Routledge.

Churchland, Paul (1981). "Eliminative materialism and the propositional Attitudes." In: *Journal of Philosophy,* 78 (February): 67-90.

Crane, Tim (2001). *Elements of Mind: An Introduction to the Philosophy of Mind.* Oxford: Oxford University Press.

Crane, Tim (2003). *The Mechanical Mind.* London: Routledge.

Crane, Tim (2013). *The Objects of Thought.* Oxford: Oxford University Press.

Crane, Tim (2014). *Aspects of Psychologism.* Cambridge, MA: Harvard University Press.

Crane, Tim (2014a). "Intentional Objects." In: Crane 2014.

Crane, Tim (2014b). "Is There a Perceptual Relation." In: Crane 2014.

Davidson, Matthew (ed.) (2007). *On Sense and Direct Reference.* Boston: McGraw Hill.

Dennett, Daniel (1987). *The Intentional Stance.* Cambridge, MA: MIT Press.

Dretske, Fred I. (1988). *Explaining Behavior: Reasons in a World of Causes.* Cambridge, MA: MIT Press.

Dretske, Fred I. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.

Dretske, Fred I. (2000). *Perception, Knowledge, and Belief: Selected Essays*. Cambridge, UK: Cambridge University Press.

Dretske, Fred I. (2000). "If You Can't Make One, You Don't Know How It Works." In: Dretske (2000).

Dummet, Michael (1973). *Frege: Philosophy of Language*. London: Duckworth.

Dummett, Michael (1991). *Les origines de la philosophie analytique.* Paris: Gallimard.

Evans, Gareth (1985). *The Varieties of Reference*. Oxford: Clarendon Press.

Everett, Anthony (2015). "*The Objects of Thought*, by Tim Crane". In: *Mind*, 124 (496):1272-1278.

Findlay, J. N. (1963). *Meinong's Theory of Objects and Values*. Oxford: Clarendon Press.

Fish, William (2010). *Philosophy of Perception: A Contemporary Introduction*. London: Routledge.

Fodor, Jerry A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.

Frege, Gottlob (1997). *The Frege Reader*. Oxford: Blackwell.

Frege, Gottlob (1997a). 'On *Sinn* and *Bedeutung*'. In: Frege 1997.

Fréchette, Guillaume (2014), (ed.) *Brentano's Concept of Intentionality: New Assessments*. Dettelbach: Röll.

Funkhouser, Eric (2006). "The determinable-determinate relation." In: *Noûs* 40 (3): 548–569.

Gallagher, Shaun, and Dan Zahavi (2012). *The Phenomenological Mind*. London: Routledge.

Ginet, Carl (2001). "Deciding to Believe." In Matthias Steup (ed.) (2011)

Goldberg, Sanford and Pessin, Andrew (eds.) (1996). *The Twin-Earth Chronicles.* New York: M.E. Sharpe.

Graham, George, Horgan, Terence, and Tienson, John (2009). "Phenomenology, Intentionality, and the Unity of mind." In Ansgar Beckermann & Brian P. McLaughlin (eds.) (2009)

Grzankowski (forthcoming). "The Determinable-Determinate Relation Can't Save Adverbialism." In: *Analysis*.

Grzankowski, Alex (2012). "Not All Attitudes are Propositional." In: *European Journal of Philosophy* (3): 374-391.

Grossmann, Reinhardt (1965). *The Structure of Mind*. Madison: University of Wisconsin Press.

Grossmann, Reinhardt (1974). *Meinong*. London: Routledge.

Grossmann, Reinhardt (1984). *Phenomenology and Existentialism: An Introduction*. London: Routledge.

Haugeland, John (1998). *Having Thought*. Cambridge, MA: Harvard University Press.

Haugeland, John (1998a). *"*The Intentionality All-Stars." In: Haugeland 1998.

Hopp, Walter (2011). *Perception and Knowledge: A Phenomenological Account*. Cambridge, UK: Cambridge University Press.

Horgan, Terence and Tienson, John (2002). "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In: Chalmers 2002.

Husserl, Edmund (2001). *Logical Investigations*. London: Routledge.

Inwagen Peter van, (2008). "McGinn on Existence." In: *Philosophical Quarterly* 58 (230): 36-58.

Inwagen, Peter Van (2009). "Being, existence, and ontological Commitment." In: David John Chalmers, David Manley & Ryan Wasserman (eds.) (2009).

Jackson, Frank (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.

Jacob, Pierre (2014). "Intentionality". In: *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2014/entries/intentionality/>.

Jeshion, Robin (ed.) (2010). *New Essays on Singular Thought*. Oxford: Oxford University Press.

Jeshion, Robin (2010a). "Singular Thought: Acquaintance, Semantic Instrumentalism, and Cognitivism." In: Jeshion 2010.

Kabasenche, William, P., O'Rourke, Michael, and Slater, Matthew H. (eds.) (2012). *Reference and Referring*. Cambridge, MA: MIT Press.

Kant, Immanuel (1974). *Kritik der Reinen Vernunft*. (2 vol.) Frankfurt am Main: Suhrkamp.

Kant, Immanuel (1929). *Critique of Pure Reason*. Trad. Norman Kemp Smith. London: MacMillan.

Kant, Immanuel (2007). *Critique of Pure Reason*. Trad. Marcus Weigelt. London: Penguin.

Kim, Jaegwon (2006). *Philosophy of mind*. Boulder, Col.: Westview Press.

Kriegel, Uriah (2007). "Intentional Inexistence and Phenomenal Intentionality." In: *Philosophical Perspectives* 21 (1): 307-340.

Kriegel, Uriah (2008). "The Dispensability of (Merely) Intentional Objects." In: *Philosophical Studies* 141 (1): 79-95.

Kriegel, Uriah (2009). *Subjective Consciousness: A Self-representational Theory*. Oxford: Oxford University Press.

Kriegel, Uriah (2011). *The sources of intentionality*. New York: Oxford University Press.

Kriegel, Uriah (2013) (ed.). *Phenomenal Intentionality*. Oxford: Oxford University Press.

Kriegel, Uriah (2015). *The Varieties of Consciousness*. New York, NY: Oxford University Press.

Kripke, Saul (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Lambert, Karel (1983). *Meinong and the Principle of Independence*. Cambridge: Cambridge University Press.

Lambert, Karel (2003). *Free Logic: Selected Essays*. Cambridge, UK: Cambridge University Press.

Larson, Richard (2002). "The grammar of intentionality." In: Georg, Peter & Gerhard, Preyer (eds.) (2002)

Levine, Joseph (1983). "Materialism and qualia: The explanatory gap." In: *Pacific Philosophical Quarterly* 64 (October): 354-61.

Lewis, David (1983). "New work for a theory of universals." In: *Australasian Journal of Philosophy* 61 (December) 343-377.

Lewis, David K. (1986). *On the Plurality of Worlds*. Malden: Blackwell.

Luntley, Michael (1999). *Contemporary Philosophy of Thought: Truth, World, Content*. Oxford: Blackwell.

Lyons, William (1995). *Approaches to Intentionality*. Oxford: Clarendon Press.

Malebranche, Nicolas (1997). *Dialogues on Metaphysics and on Religion*. Cambridge: Cambridge University Press.

McCulloch, Gregory (2003). *The Life of the Mind: An Essay on Phenomenological Externalism*. London: Routledge.

McDaniel, Kris (2009). "Ways of being." In: David John Chalmers, David Manley & Ryan Wasserman (eds.) (2009)

McDowell, John (1998). *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

McDowell, John (2009) *Having the world in view: essays on Kant, Hegel, and Sellars*. Cambridge, MA: Harvard University Press.

McGinn, Colin (1991). *The Problem of Consciousness*. Malden: Blackwell.

McGinn, Colin (2000). *Logical Properties*. Oxford: Clarendon Press.

Mackie, J. L (1982). *The Miracle of Theism*. Oxford: Clarendon Press.

Marcus, Ruth Barcan (1961). "Modalities and Intensional Languages." In*: Synthese, 13(4): 303–322.*

Meinong, Alexius (1915). *Über Möglichkeit und Wahrscheinlichkeit*. Leipzig: Ambrosius Barth Verlag.

Meinong, Alexius (1960). "On the theory of objects." In: Chisholm, Roderick (ed.), *Realism and the Background of Phenomenology*. Free Press.

Mendelsohn, Richard L. (2005). *The Philosophy of Gottlob Frege*. Cambridge: Cambridge University Press.

Mill, J.S. (1974). *A System of Logic*. Toronto: Toronto University Press.

Moltmann, Friederike (2013). "The semantics of existence." In: *Linguistics and Philosophy* 36 (1):31-63.

Montague, Michelle (2007). "Against propositionalism." In: *Noûs*, 41 (3): 503–518.

Montague, Michelle (2011). "The Phenomenology of Particularity." In: Bayne and Montague 2011.

Montague, Michelle (2015). "Review of: *The Objects of Thought*, by Tim Crane". In: *Analysis,* Vol 75, Nb. 2.

Moore, Adrian, W. (2006). *Points of View*. Oxford: Clarendon Press.

Moran, Dermont (1996). "Brentano's Thesis." In: *Proceedings of the Aristotelian Society* 70 (70):1-27.

Moran, Dermot (2000). *Introduction to Phenomenology*. London: Routledge.

Nagel, Thomas (1974). "What Is It Like to Be a Bat?" In: *The Philosophical Review*, 83 (4): 435–450

Nagel, Thomas (1986). *The View from Nowhere*. New York: Oxford University Press.

Parsons, Terence (1980). *Nonexistent Objects*. New Haven: Yale University Press.

Perszyk, Kenneth J. (1993). *Nonexistent Objects*. Dordrecht: Springer Verlag.

Peter, Georg & Preyer, Gerhard, (eds.) (2002). *Logical Form and Language*. Oxford: Oxford University Press.

Pitt, David (2012). "Mental Representation." In: *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), URL=<https://plato.stanford.edu/archives/fall2013/entries/mental-representation/>.

Priest, Graham (2005). *Towards Non-being*. Oxford: Oxford University Press.

Priest, Graham (2006). *In Contradiction*. Oxford: Oxford University Press.

Priest, Graham (2008). *An Introduction to Non-classical Logic*. Cambridge: Cambridge University Press.

Prior, A. N. (1971) *Objects of Thought*. Oxford: Clarendon Press.

Putnam, Hilary (1981). *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Quine, W. V. (1953). *From a Logical Point of View*. Cambridge, MA: Harvard University Press.

Quine, W. V. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Recanati, François (1993). *Direct Reference*. Oxford: Blackwell.

Recanati, François (2010). "In Defence of Acquaintance." In: Jeshion (ed.) (2010).

Recanati, François (2012). *Mental Files*. Oxford: Oxford University Press.

Reimer, Marga & Bezuidenhout, Anne (ed.) (2004). *Descriptions and Beyond*. Oxford: Oxford University Press.

Routley, Richard (1980). *Exploring Meinong's Jungle and Beyond*. Canberra: Research School of Social Sciences, Australian National University.

Rosefeldt, Tobias (2011). "Frege, Pünjer, and Kant on Existence." In*: Grazer Philosophische Studien,* 82(1) 329–351.

Rosefeldt, Tobias (Unpublished). *Was es nicht gibt. Eine Untersuchung des Begriffes der Existenz*. Habilitationsschrift an der Philosophischen Fakultät der Ruprecht-Karls-Universität Heidelberg 2006

Russell, Bertrand (1905). "On Denoting." In: *Mind,* New Series, Vol. 14, No. 56.

Russell, Bertrand (1911). "Knowledge by Acquaintance and Knowledge by Description." In: *Proceedings of the Aristotelian Society* 11.

Russell, Bertrand (1919). *Introduction to Mathematical Philosophy*. London: George Allen and Unwin.

Russell, Bertrand (2010). *The Philosophy of Logical Atomism*. London: Routledge.

Sainsbury, R.M. (2004). "Referring Descriptions." In: Reimer & Bezuidenhout (ed.) (2004)

Sainsbury, R. M. (2005). *Reference without Referents*. Oxford: Clarendon Press.

Sainsbury, R. M. (2010). *Fiction and Fictionalism*. London: Routledge.

Sainsbury, R. M. (2010). "Intentionality without exotica." In Robin Jeshion (ed.) (2010)

Sakar, Sahotra & Pfeifer, Jessica (eds.) (2006). *The Philosophy of Science. An Encyclopedia*. London: Routledge.

Salmon, Nathan (1986). *Frege's Puzzle*. Cambridge, MA: MIT Press.

Schiffer, Stephen R. (2003). *The Things We Mean*. Oxford: Clarendon Press.

Searle, John R. (1969). *Speech Acts*. Cambridge: Cambridge University Press.

Searle, John R. (1983). *Intentionality.* Cambridge: Cambridge University Press.

Searle, John R. (1992). *The Rediscovery of the Mind.* Cambridge, MA: MIT Press.

Siewert, Charles (1998). *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.

Siewert, Charles (2011). "Phenomenal thought." In: Bayne, Tim and Montague, Michelle (ed.) (2011)

Soldati, Gianfranco (2008). "Representation, Intentionality, and Consciousness." In: Holderegger, Adrian; Rager, Günter & Sitter-Liver, Beat (eds.) (2008): *Hirnforschung und Menschenbild.* Basel: Schwabe.

Strawson, Galen (1994). *Mental Reality.* Cambridge, MA: MIT Press.

Strawson, Galen (2008). *Real Materialism and Other Essays.* Oxford: Clarendon Press.

Strawson, Galen (2011). "Cognitive phenomenology: real life." In: Bayne, Tim & Montague, Michelle (eds.) (2010)

Steup, Matthias (ed.) (2001). *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.

Tahko, Tuomas E. (2015). *An Introduction to Metametaphysics.* Cambridge, UK: Cambridge University Press.

Tarski, Alfred (1965). *Introduction to Logic and to the Methodology of Deductive Sciences.* New York: Oxford University Press.

Tienson, John *(2013).* "Kasimir Twardowski on the content of presentations." In: *Phenomenology and the Cognitive Sciences 12 (3): 485-499.*

Thomasson, Amie L. (1999). *Fiction and metaphysics.* Cambridge, UK: Cambridge University Press.

Twardowski, Kasimir (1977). *On the Content and Object of Presentations.* The Hague: M. Nijhoff.

Valberg, Jerome J. (1992). *The Puzzle of Experience.* Oxford: Oxford University Press.

Voltolini, Alberto (2006). *How Ficta Follow Fiction: A Syncretistic Account of Fictional Entities.* Dordrecht: Springer.

Walters, Lee (2015). "The Problem of Non-Existence: Truth-Making or
    Semantics?" In: *Disputation*, Vo. VII, No. 41, November 2015.
Zalta, Edward N. (1988). *Intensional Logic and Metaphysics of Intentionality*.
    Cambrige (MA): MIT Press.