

An Agent of Attention: An Inquiry into the Source of Our Control

by

Aaron George Bruce Henry

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy Graduate Department
University of Toronto

© Copyright Aaron George Bruce Henry 2019

An Agent of Attention: An Inquiry into the Source of Our Control

Aaron George Bruce Henry

Doctor of Philosophy

Graduate Department of Philosophy
University of Toronto

2019

Abstract

When performing a skilled action—whether something impressive like a double somersault or something mundane like reaching for a glass of water—you exercise control over your bodily movements. Specifically, you guide their course. In what does that control consist? In this dissertation, I argue that it consists in attending to what you are doing.

More specifically, in attending, agents harness their perceptual and perceptuomotor states directly and practically in service of their goals and, in doing so, settle the fine-grained manner in which their bodies will move—details an intention alone leaves unsettled. This requires, among other things, that we reject views on which agents' control is identical with their practical rationality.

When all goes well, agents attentionally prioritize what is motivationally relevant to them to the exclusion of what would otherwise distract them from achieving their goals. However, sometimes agents attend *distractedly*—i.e., *without* prioritizing. As the aim of

attention is to avoid distraction, this entails the possibility of *defective* attention. Defective attention, in turn, casts light on scenarios in which agents lose control over what they are doing, as when a skilled practitioner ‘chokes under pressure’.

A complaint sometimes levelled against accounts, like mine, that claim to reduce agents’ control of their behaviour to that of causally efficacious mental states or events is that these accounts invariably deprive *agents themselves* of their rightful role in the generation of behaviour. This is the “Disappearing Agent Problem” for “reductive” or “event-causal” theories of action. I argue that, correctly understood, extant reductive theories *do* face a genuine Disappearing Agent Problem. However, it is a problem we solve by recognizing the role that conscious attention plays in making an action the agent’s own. Accordingly, I develop and defend an attentional account of *action ownership*. On this view, allocating conscious attention in service of your goals is sufficient for a kind of conscious perspective (“motivational perspective”), which, when active in controlling your behaviour, constitutes the behaviour as your own doing. As I explain, such perspective also contributes to explaining the subjective structure of an agent’s perceptual awareness of the world around her.

Acknowledgments

My primary debt is to my supervisor, Imogen Dickie. Imogen has been unwavering in her support for my research and her dedication to my development as a scholar. I could not have hoped for a more philosophically serious and exacting supervisor. I hope to live up to her example as a philosopher and teacher.

I have also greatly benefitted from the incisive and generous feedback of my other committee members, Mohan Matthen and Diana Raffman, and from the probing commentaries of my two examiners, Chris Mole and Bill Seager. I want to thank all of them for their warm encouragement about my project and for the care and hard work they have put into critiquing it.

In addition to my committee, I want to express my gratitude to other faculty members, both at the University of Oxford and the University of Toronto, who have been generous with their time and, at various stages, contributed to the formation and trajectory of my project. While I was a BPhil student at Oxford, Tim Bayne and Martin Davies both encouraged me to work on attention and oversaw my initial work on the relationship between attention and consciousness. Although the philosophical problems that I engage with here are quite different from those that I was working on then, traces of that early work are nevertheless detectable. I thank Tim and Martin for this early support. Later, at Toronto, I benefitted from conversations with many of the department's wonderful faculty. Early conversations and collaborative work with Evan Thompson on the sense of self informs some of my central claims about the sense of agency in Chapter 3. A reading group that I did with Phil Clark in the spring of 2016 was especially important for clarifying the action-theoretic issue I had been grappling with in Chapter 1. I have also

profited from discussing Leibniz's views about the sense of agency with Marleen Rozemond and the role of thought in skilled action with Sonia Sedivy.

At Toronto, I have also been fortunate to be among a group of highly talented and intellectually generous graduate colleagues. Without meaning to be exhaustive, I would like to express my thanks to: Dominic Alford-Duguid, John Bunke, Elliot Carter, Kevin Connolly, Alex Djedovic, Lisa Doerksen, Zac Irving, Damian Melamedoff, Parisa Moosavi, Jonathan Payton, Adrienne Prettyman, Bryan Reece, Luke Roelofs, Amogh Sahu, Sean Smith, Dave Suarez, and Mason Westfall. I'm particularly grateful to Zac Irving, with whom I've spent countless hours debating and exploring ideas at the heart of this dissertation at the early stage of writing, and to Amogh Sahu, with whom I had many valuable conversations at the later stages.

I have been fortunate enough to receive many excellent commentaries from the broader philosophical community about various components of my project. For their valuable conference commentaries, I thank: Carolyn Jennings, Alex King, Sarah Robins, Jonathan Simon, Daniel Stermer, Jody Tomchisen, Robyn Repko Waller, and Wayne Wu. I am also grateful to have had constructive dialogue at workshops, conferences, and other venues with: Kristen Andrews, Paul Bello, Michael Brent, Joe Levine, Myrto Mylopoulos, Miguel Angel Sebastian, Josh Shepherd, Neil Sinhababu, John Schwenkler, Matthew Soteriou, Helen Steward, and Sebastian Watzl. The generous feedback I've received from these philosophers and scientists has been extremely helpful.

Finally, I thank my friends, family, and partner, Anastasia. The process of writing a dissertation was sometimes a challenging and lonely experience. The support of my loved ones made it much less so.

The research for this dissertation was supported by Social Science and Humanities Research Council's Joseph-Armand Bombardier Canada Graduate Scholarship.

Contents

1	Attention in Action	1
1.1	Introduction	1
1.2	The motor control challenge	2
1.3	A solution to the motor control challenge	6
1.4	The extent of agential control: Rebutting the intellectualist	10
1.4.1	Motor control as rational	11
1.4.2	Motor control as “subpersonal”	15
1.4.2.1	Motor control as unconscious	15
1.4.2.2	Motor control as motivationally unguided	18
1.4.2.3	Motor control <i>is</i> motivationally guided	20
1.5	Coping with distraction	23
1.6	Comparison with extant accounts	27
1.6.1	Cognitive unison theory	27
1.6.2	The “could have done otherwise” model	29
2	Actionism and Attention Capture	36
2.1	Introduction	36
2.2	The distinction between voluntary and automatic attention	39
2.3	The environmental salience view	42
2.4	Against the environmental salience view	46
2.5	The priority view of automatic attention	51
2.5.1	Agential priorities, acts of prioritization, and the integrated priority map	52
2.5.2	Failures of prioritization	57
2.5.3	An alternative to “rigorism” and “permissivism”	64
2.6	Conclusion	66
3	Attention and the Problem of the Disappearing Agent	71
3.1	Introduction	71

3.2	“The Disappearing Agent Problem”	73
3.2.1	Ambitious formulations	74
3.2.2	Velleman’s Problem (the Shrinking Agent Problem)	76
3.2.3	Already owned?	79
3.2.4	Yes, this is agency (not <i>merely</i> sense of agency)	83
3.3	A puzzle about the sense of ownership in “immersive” action	84
3.4	The perspectival structure of an experience of acting is based in the causal structure of conscious attention	91
3.5	Extension to non-immersive experiences of acting	98
3.6	Losing ownership?	100
4	Attentional Structure	108
4.1	Introduction	108
4.2	The sharpening and boosting of perceptual appearances	111
4.3	The replication argument	117
4.3.1	Replying to the replication argument: Step 1	117
4.3.2	Replying to the replication argument: Step 2	121
4.4	Modulating perceptual organization with attention	126
4.5	What appearance replicas show: the inalienability of conscious attention	130
4.6	Conclusion	135

Chapter 1

Attention in Action

1.1 Introduction

What are the features of behaviour that qualify it as manifesting the agent's control? A traditional "intellectualist" answer identifies agential control with practical rationality and exercises of agential control with processes like practical reasoning culminating in decision. However, intellectualists face a challenge in explaining the control agents exercise when moving their bodies. To illustrate, when an expert squash player performs a winning drop shot, the control she manifests does not terminate with her decision, e.g., to perform a drop rather than a boast. Intuitively, she also manifests control in the precise manner in which she executes her decision—e.g., in such fine details as with what grip to hold the racquet, how high to lift her arm, and how far and with what force to follow through. Assuming that agents do not explicitly deliberate and decide on these fine details of their movements, intellectualism leaves us without an adequate understanding of how agents settle such properties of their actions. This is the *motor control challenge* to intellectualism.

The motor control challenge really is fatal for intellectualism. We must extend our attributions of agential control beyond deployments of the agent's rationality to include certain of the psychological states and processes used in fine-grained intention execution. In rejecting intellectualism as an adequate theory of agential control, however, we confront a new explanatory demand. We must say what an agent's control of her fine-grained movements consists in psychologically if not rational deployments of the sort intellectualists have traditionally appealed to in their accounts of agency. Moreover, an adequate account of these psychological states and processes must simultaneously reveal why such states and processes qualify as deployments of the *agent's* control.

To see why the latter may be an issue, consider the plausible suggestion that, in accounting for how an agent's fine-grained bodily movements are controlled, we must

appeal to finer grained states than beliefs, desires, and intentions, such as the agent's perceptual states. Unless agents somehow base their actions on perception, an obvious feature of both ordinary and highly skilled actions would remain mysterious—namely, how agents tailor their actions flexibly and precisely to particular features of a situation, while respecting strict temporal constraints on intention execution. A role for perception in accounting for fine-grained human motor control thus seems unavoidable. On the other hand, unless we explain how these action-informing perceptual states become anchored to the agent's goals in acting, we will not have explained why perceptually based behaviours qualify as manifesting the *agent's* control. It thus seems that understanding perception's central role in enabling skilled, intentional action requires an account of *motivational harnessing*: the process whereby perceptual states become coordinated with the agent's goals in acting and, as a result, structured according to the standards of success that the latter set for behaviour.

I'll argue that *attending* just *is* this process. In attending, an agent's perceptual states become suitable to generate intention-congruent motor representations and thereby to inform fine-grained bodily movement. The process of attending is the process whereby agents allocate control toward a motivational state's fine-grained motor implementation.

The plan for the chapter is this. In §1.2, I sketch the motor control challenge to intellectualism and lay some initial groundwork for a non-intellectualist account of agential motor control. In §1.3, I present my own proposal in terms of attending. In §1.4, I return to the intellectualist and argue that their main lines of response are unsuccessful. In §1.5, I consolidate my proposal by explaining the role attention plays in dealing with distractions. Finally, in §1.6, I compare my account with others that emphasize a role for attention in action. Appreciation of the difficulties facing these proposals will strengthen the case for my own conception of attention and its role in action.

1.2 The motor control challenge

According to what I'll call an "intellectualist" conception, the control agents manifest in acting can be characterized exhaustively in terms of two factors. The first is the causal relation between behaviour and the agent's propositional attitudes. The second is psychological processes operating on these attitudes (e.g., practical reasoning). Davidson's (1963) theory of action offers an example. According to Davidson, an action is an event that is intentional, and an event is intentional if and only if (henceforth, "iff") it is caused by reasons that motivate and rationalize it—on Davidson's original view, a desire to Φ and a belief that Ψ -ing is a way to Φ . To illustrate, suppose Athena visits a foreign city and wishes to see City Hall, believes that taking the B-32 bus is a way to get there and

that taking the B-32 is superior way of going there than her available alternatives (e.g., taking a taxi, walking, etc.). Other things being equal, she will take the B-32, and her doing so will have been intentional iff caused (non-deviantly) by her motivating complex of beliefs and desires—i.e., her reasons.¹ On this view, behaviour is attributable to the agent rather than to an external event or to one of the agent’s subsystems iff it manifests the agent’s *rationality*. It is this alignment of agential control with rational control that qualifies Davidson’s account as “intellectualist” in my sense.

Some philosophers have objected that Davidson’s account fails to do adequate justice to the role of the *agent* in action because his account does not do adequate justice to the role of the agent’s *rationality* in action. This is then taken to motivate the introduction of yet other intellectual requirements on action. For example, it has been suggested that that a full-fledged action must not only be motivated by reasons, but by reasons that the agent “endorses” and “reflectively identifies with” (e.g., Velleman 1992; 2000; Bratman 1996). The objection with which I am concerned in this chapter goes in the opposite direction. According to this objection, Davidson’s view fails to adequately accommodate the agent’s role in action because it does not do adequate justice to a different kind of control agents exercise in acting. Specifically, the account does not accommodate the control agents exercise in skillfully moving their bodies.

To see the problem, consider the act of taking a sip of water. On Davidson’s view, we can understand this as follows. If you are thirsty and believe that taking a sip from the glass of water in front of you is a way to relieve your thirst, then, all else equal, you will take a sip from the glass before you, and your sip will have been intentional iff non-deviantly caused by your reason for action. However, as Israel, Perry, and Tutiya (1993) note, Davidson’s account leaves open a “gap” between the propositional contents of the agent’s motivating belief-desire complex—which in our example concerns a distal object upon which the agent intends to act (the glass of water)—and the local bodily movements required to act on the object in the intended way. To bring out the difficulty, Israel et al. discuss the “wrong movement problem”. This is the problem that some action failures seem not to be due to any error in one’s orienting beliefs—e.g., beliefs about one’s circumstances or about which actions are ways of accomplishing which ends in the circumstances—but to failures of *execution*. For example, one might fail to take a sip from the glass in front of oneself not because one falsely believes there is a glass of water located there or that sipping from it is a way to quench one’s thirst, but because one executes the wrong movements when reaching for the glass. In such a case, action failure doesn’t seem to arise from any falsity in the agent’s beliefs, at least as these are standardly

¹Davidson later came to hold that intentions do not reduce to belief-desire pairs and instead are a *sui generis* attitude (Davidson 1978). He “dispairs” about deviant causal chains in Davidson (1973).

conceived. For example, one's beliefs that there is a glass there and that reaching for it is a way to satisfy one's thirst are true and, let's suppose, justified. Rather, the action failure seems to arise from lack of practical intelligence or "know how". To highlight the central role that the capacity for motor control plays in this challenge to Davidson, I'll call it the "motor control challenge".

Israel et al.'s own response to the motor control challenge is not to reject Davidson's intellectualism. They instead respond by enriching our account of the propositional attitudes on the basis of which agents act to include beliefs about *how* to execute different kinds of intentional actions ("belief how") (Israel et al. 1993, 534). With these beliefs, Israel et al. aim to bridge the gap that they claim to have found in Davidson's account between distal objects and local bodily movements. However, as Pacherie (2011, 69) notes, while Israel et al.'s proposal might explain some action failures, it cannot explain all. What it seems able to explain are execution failures resulting from false beliefs about which bodily movements are ways to execute a given sort of action. But it seems not to explain what goes wrong when the agent correctly believes that executing a certain type of movement in the circumstances is a way to grasp the glass (and as a result intends to execute that type of movement) but fails to execute her intention because her body does not move as she intended. By hypothesis, failures of the latter sort don't result from any fault in the agent's beliefs or other attitudes, but from a failure to *implement* them. Assuming the latter execution failures can reflect a lack of agential control, then it seems agential control cannot be the same as rational control.

The motor control challenge forces the intellectualist either to explain why the above action failures, contrary to initial appearances, are rational failures or why these action failures, contrary to initial appearances, are not failures of agential control. In §1.4, I'll assess these potential ways of meeting the challenge. For the moment, I want to point out that the motor control challenge also raises an explanatory demand on those raising the objection to the intellectualist. In particular, they must say what an agent's motor control might consist in psychologically, if not deployments of rational processes of the sort intellectualists have traditionally invoked in their accounts of agency. This is the explanatory demand I want to take up now. Because Pacherie (2006; 2008; 2011) herself offers the first step toward an alternative account with the notion of a "nonconceptual motor representation", and because my account will build on this notion, I'll start with it.

The notion "motor representation" is taken from motor control theory (see e.g., Jeannerod 1997; 2006). It refers to the type of representation that is the proximate cause of intentional movement. These representations are also responsible for ensuring the satisfaction of strict biomechanical, kinematic, and temporal constraints on the successful implementation of an intention. In order to perform these roles, they must specify the

precise movement parameters for an action (e.g., with what force to move one's arm) and they must specify an action's precise outcome (e.g., the specific grasp that will obtain when the action is complete).

Unlike intentions, motor representations are nonconceptual (Pacherie 2011). A representational state is "nonconceptual" if its instantiation doesn't require its subject to possess the concepts needed to characterize its content (e.g., to express its correctness conditions), where a "concept" is a representation deployable in thought (Peacocke 1992). One reason for regarding motor representations as nonconceptual rather than conceptual concerns the fineness of grain or determinacy of their content compared with that of intentions.² For example, in order to successfully reach for a glass of water and take a sip, exact numerical magnitudes of various movement parameters must be represented at some level. To claim that intentions are responsible for this is to claim that we have intentions that represent, at the conceptual level, the exact numerical magnitudes for such properties as the speed, trajectory, force, and final location of one's reach or the aperture and force of one's grip. This claim is implausible. Given the capacity limitations of conceptual states and processes—e.g., the coarseness of their content and the slowness of explicit reasoning—it's unlikely the agent represents the required information conceptually in the form of beliefs and intentions. It seems much more likely that the representational states responsible for specifying these movement parameters and outcomes are nonconceptual motor representations.

As noted earlier, nonconceptual motor representations provide the beginning of an account of the psychological basis of agential motor control. A complete account, however, would explain why we should think of these psychological states as contributing a distinctively *agential* form of control. My own answer builds on a feature of Davidson's view that, I will suggest, survives the motor control challenge to his intellectualism:

MOTIVATIONAL CONTROL: An agent exercises control over some property of her behaviour iff that property of her behaviour is the result of control by a motivational state of hers—e.g., by the intention with which she acts.

What the motor control challenge ultimately shows, in my view, is that intellectualists like Davidson have operated with an overly restrictive view of the mental states and processes constitutive of exercises of motivational control (cf. Dickie 2015, 95 ff). On their view, the way one settles the answers to questions about how one will act (e.g., whether or not one will Ψ) is by engaging in explicit inference. This process culminates in the rational selection of a specific course of action—e.g., a conceptual intention to Ψ given

²Evans (1982, 229) defends the claim that perception possesses nonconceptual content on similar grounds.

one's goal of Φ -ing. In deploying such rational processes, a distinctively rational form of control is thereby exercised. More specifically, a distinctively *practical* form of rationality is exercised. This is a form of inference in which one adopts a conclusion on the basis of certain premises, not because the conclusion is "shewn to be true by the premises" (Anscombe 1957/2000, 58), but rather because its premises show it to be good in light of one's goals or to count in favour of the conclusion. Supposing agents exercise not only rational control but motor control, we can ask about the process or processes that agents use to settle finer grained questions about how they will implement their goals—e.g., whether or not they will Φ in determinate way w . Given MOTIVATIONAL CONTROL, this would be a psychological process that is anchored to the agent's intention. But rather than resulting in a further implementational intention (e.g., to Ψ as a means of Φ -ing), this process instead results in a nonconceptual motor representation. The challenge facing the opponent of intellectualism is to describe such an intention-directed psychological process without relapsing into intellectualism.

In the next section, I sketch an account of the psychological process whereby an agent's intention to act generates a nonconceptual motor representation and, in doing so, settles the answers to finer grained questions of intention implementation than an intention alone can settle.

1.3 A solution to the motor control challenge

In this section, I propose that it is through attention that agents settle the fine details of their actions and thereby exercise motor control. In particular, attending is the process of harnessing perceptual information to coherently resolve implementational questions that remain unresolved by the agent's conceptual intention to act. I'll start by explaining how I conceive of attention as a psychological phenomenon. I'll then explain how the features of this causal process help us to answer the motor control challenge.

As I shall understand it, attending is a biasing process that makes a subject differentially prone to use some available information at the expense of other information. To illustrate, suppose you begin to look for your favourite red pen with the aim of writing something down. Your intention to find your red pen introduces a bias into visual processing, so that you become preferentially sensitive to red, pen-shaped stimuli to the exclusion of other stimuli. Once it is detected, motor programming for reaching commences. You reach for your pen with a grip appropriate for writing, and you consequently do as you intend.

I'll now say more about how I propose to understand this biasing process.

We can understand the process of attending as comprising two constituent mental

states: what I'll call an "attentional source state" and an "attentional resultant state". An "attentional source state" is the mental state that drives the bias, whereas an "attentional resultant state" is the effect of this bias on states of sensorimotor processing. One virtue of this conception of attention is that it helps to resolve an ambiguity that one sometimes finds in the empirical literature regarding where in this biasing process attentional selection *itself* exists or occurs. I'll illustrate this with the case of visual attention. According to more traditional views of visual attention, the realizer of attention itself is identified with the causal source of attentional bias in a dedicated "supramodal", nonsensory area of the brain that acts on lower level visual systems (e.g., Posner and Peterson 1990). In my terms, this approach identifies attention with the attentional source state. By contrast, more contemporary models often identify visual attention with the resultant (or "emergent") effect of some biasing process *on* states of visual processing (e.g., Desimone and Duncan 1995; Duncan 1998). These generate opposing conceptions of attention as either essentially cognitive and high level or else sensory and low level (see Stinson 2009; Allport 2011; Mole 2015; Wu 2017 for discussion). My conception has elements of each picture. On my conception, a *process* of visual attention is realized in the *biasing* of lower level visual systems *by* higher level source states, and a *state* of visual attention is realized in the modulated state of visual processing that results from this biasing process. Consistent with contemporary models, a state of visual attention is, on this conception, a sort of visual state. However, there is also a constitutive role for participating source states in making that visual state one of *attention*. That is, a state of visual attention is a visual state whose status as attention constitutively depends on the biasing influence of certain source states. I return to this framework often throughout the dissertation.

When the source state for an attentional episode is the subject's goal or intention, the attention is said to be "goal-directed" or "top-down" as opposed to "stimulus-driven" or "bottom-up" attention. In the next chapter, I argue that this traditional dichotomy between two types of attentional control—one intuitively active; the other intuitively passive—are less clear-cut than sometimes assumed. I will argue there that an agent's attention is always partly driven by a motivational source, and so in an important sense is always motivated or "goal-directed". In this chapter, however, my concern is specifically with attention as it is deployed in service of an agent's intention to achieve some goal. Consequently, I am concerned with an unambiguously goal-directed form of attention. I will postpone consideration of the different varieties of attentional control—i.e., the different potential source states of attentional bias—till the next chapter. For the remainder of this chapter, unless otherwise stated all references to "attention" should be read as referring to goal-directed attention—in particular, to attention that is directed by the agent's intention to accomplish some goal.

The attentional source state for goal-directed attention can be thought of a motivational state, such as an intention. With reference to the tasks given to subjects in laboratory conditions, psychologists refer to a subject's "intentional set", "task set", or "attentional set" as playing this source role. We can think of a task set as the intention that sets the goal for psychological processes and directs them toward that goal's achievement. Thus, in adopting a certain task set, the subject's intention to complete that task comes to serve as an attentional source state. In our earlier example, your intention to write with your favourite red pen serves as your intentional set, which introduces a top-down bias into the network favouring task-relevant stimuli and responses and suppressing or inhibiting task-irrelevant stimuli ("distractors") and responses. Intuitively, in adopting this task set, you are giving ("setting") yourself the task and, in so doing, becoming mentally prepared ("set") for the task ahead. This means anchoring your attention to that intention.

If the process of top-down biasing by intention disposes a subject to respond to a target in a task-congruent manner, does this show that attending operates not only on the subject's perceptual states but also on her motor states? The answer depends on how we draw the boundary between perceptual and motor states, and the framework I'm proposing is neutral about how we draw this line. One option is to distinguish states of perceptual attention and motor attention or perceptual and motor aspects of a state of attention. This would accord with the distinction sometimes drawn between "selection-for-perception" and "selection-for-spatial-motor action". According to this framework, two forms of attention, subserved by partially functionally and anatomically separate processing streams, become coordinated under a single "visual attention system" (Schneider and Deubel 2002; cf. Cisek 2007; Memelink and Hommel 2013; Herwig 2015 for complementary approaches). If by "perception", however, one includes not only the ventral visual states used in categorizing and identifying objects for uptake in reasoning and planning, but also the visual states immediately used in fine-grained movement-control within the dorsal visual processing stream (Milner and Goodale 1995/2006), then it is possible that parameter-setting for a motor representation can be fully specified once perceptual (or "perceptuomotor") states have been attentionally modulated. The latter visuomotor states are sometimes said to represent objects in terms of the fine-grained movements required to, for example, reach and grasp an object. If we include the latter as perceptual, then perceptual attention may be sufficient to resolve determinate movement details because, among the perceptual states that receive attentional biasing, some represent the agent's environment in terms of the determinate movements required to perform various object-directed actions, and this may be sufficient to generate a motor representation. In either case, the subject's intention biases psychological processing within lower level, nonconceptual systems in ways that eventually culminate in formation of a motor

representation.

At the end of the last section, I formulated the motor control challenge for the opponent of intellectualism as follows: to accommodate the claim that in acting, an agent's bodily movements are motivationally controlled, while resisting the claim that motivational control reduces to conceptual decision-making. I suggested that this demand could be met by identifying a process whereby agents select a means to an intended end that does not require that the selected means figure in the conceptual content of a decision. I now want to suggest that such a nonconceptual form of selection occurs when agents attend in acting. By attending, the agent selects solutions to implementation problems in a way that is simultaneously: (i) directed by the agent's intention, and (ii) distinct from decision-making.

That attending is an intention-directed process follows from an intention serving as its source state. The primary difference between attending and practical reasoning is that, in attending, the mental states that one's intention harnesses in order to settle its own implementation are not conceptual states, such as orienting beliefs about one's circumstances and which ways of acting are conducive to achieving one's goals in those circumstances. Instead, in attending, one's intention harnesses non-conceptual states. Nevertheless, important parallels exist between attending and practical reasoning. Just as practical reasoning is a process of harnessing beliefs in order to settle high level questions about how one will act (e.g., whether one will Ψ given one's goal of Φ -ing), attending is a process of harnessing nonconceptual information in order to settle finer grained questions about how one will act (e.g., the specific manner in which one will Ψ given one's goal of Φ -ing). Just as agents put their beliefs to use as premises in deliberating about how to act, they put their nonconceptual perceptual states to use in attending in order to settle how they will move their bodies. The former process leads ultimately to a conceptual action-guiding state (an intention); the latter to a non-conceptual action-guiding state (a motor representation). When acting with skill, agents exploit, in a practical form, their status as both reasoners and as perceivers.

Related to the non-conceptual status of perceptual attention, we can identify two senses in which the selection the agent makes by attending will typically be "automatic" rather than deliberate or intended. First, I've said that the final product of an agent's attentional selection is a nonconceptual motor representation responsible for executing the fine-grained movement parameters of an intentional action. The content of this representation will include, among other things, precise magnitudes like velocity and grip aperture. Since intention, by contrast, is a conceptual state, the fine-grained properties the motor representation specifies need not, and typically will not, figure in the content of the agent's explicit intention. As a result, when a motor representation successfully pro-

duces the movement that its content specifies, the properties of the movement for which the motor representation is responsible need not be intended by the agent. For example, when you finally reached for your red pen in order to write something down, you likely did not deliberate about the precise grip size and aperture that you used to grasp it. Instead, such fine-grained properties of your action were performed “automatically” or “sub-intentionally”. Relatedly, while subjects can usually offer some verbal explanation for what they have done intentionally, they may not be in a position to do the same for properties of an action that are the product of nonconceptual motor representations. For example, if asked why you just grabbed your red pen, you might respond “Because I needed to write something down with it”. But if asked to justify some highly specific property of your reach, you may not be able to offer the same kind of justification because you did not realize your grip had the described property (cf. Luthra 2016, 2273). The relevant property was an automatic rather than deliberate feature of your action.

A second dimension of automaticity comes from the fact that attending itself will often be automatic in that agents typically will not intend to attend as a means to achieving their goals. For example, having formed the intention to write something down with your pen, you do not typically then strategically decide to focus your attention on the desk in order to find your pen (though you can do this). Rather, once an intention is formed, it normally comes to structure attention automatically. Because attention is deployed automatically, subjects can very well be surprised to learn how they were directing their visual attention to a scene while performing a task (cf. Wu 2014b, 35). Rather than being initiated on the basis of practical reasoning, attending typically occurs in tandem with one’s practical reasoning. Attending is a motivated, yet typically automatic mode of agential control.

In this section, I’ve sketched my proposed solution to the motor control challenge. Attending is the process whereby agents settle finer grained properties of intention implementation than are normally settled by decision. This process is realized in the biasing by an agent’s intention of nonconceptual perceptual and perceptuomotor states. In so doing, we exercise automatic, yet motivationally directed control over our bodily movements. Having sketched my proposal, I want now to return to the intellectualist and explain why we should prefer my account over theirs.

1.4 The extent of agential control: Rebutting the intellectualist

As I noted in §1.2, intellectualists about agential control face two options in responding to the motor control challenge. First, they could accept that the motor control manifested in

fine-grained behaviour is a genuine exercise of agential control but argue that this form of control can be accommodated by suitably enriching our conception of the propositional attitudes on the basis of which agents act. Second, they could grant that motor control cannot be fully explained as a form of rational control, but then deny that they are under any burden to explain motor control because motor control doesn't instantiate a genuinely agential form of control. The former response claims to bring motor control within the scope of agency, while the second claims to exclude it. I'll consider each of these responses in turn, and argue that neither provides a convincing response for the intellectualist. Consequently, we should reject intellectualism.

1.4.1 Motor control as rational

According to the first reply I'll consider, the fine-grained properties of action reflect the agent's rational control. In particular, on this proposal, an agent who lacks the ability to successfully execute a certain type of intentional movement does not possess the same beliefs as a person who possesses that ability.

One question for someone pursuing this response is whether or not they regard the proposition believed by the skilled agent as being sufficiently fine-grained to explain the precise movements the agent makes in a particular situation, at a time, in response to specific cues, etc. There are two worries facing an affirmative answer to this question. The first is that it will place unrealistic demands on decision-making. When one intentionally reaches for a glass of water or executes some skilled performance, it is true that precise magnitudes governing various movement parameters of the reach must be represented at some level. But it is implausible that such information is represented in the agent's decision, presumably on the basis of deliberation, about how to execute the action. Second, even supposing that an agent possessed beliefs that were fine-grained enough to rationalize the precise execution details of her action (and that all of these beliefs were true and justified), this still wouldn't help to explain the case where the execution error arises in the *application* of those beliefs to a particular situation.

To circumvent these difficulties, the intellectualist might appeal to Stanley and Williamson's (2001), Stanley's (2011), and Stanley and Krakauer's (2013) account of "know how" as a species of propositional knowledge and of skilled action as an action governed by such knowledge. Stanley and Williamson illustrate their proposal with Hannah's knowledge of how to ride a bicycle:

Hannah knows [PRO how to ride a bicycle] is true relative to a context *c* if and only if there is some contextually relevant way *w* such that Hannah stands in the knowledge-that relation to the Russellian proposition that *w* is a way

for Hannah to ride a bicycle, and Hannah entertains this proposition under a practical mode of presentation. (Stanley and Williamson 2001, 430)

Importantly, the key element in this proposal is not the proposition Hannah knows—i.e., that w is a way for Hannah to ride a bicycle. In general, nonexperts can know many of the same propositions experts know about which actions are required to execute a certain skill—e.g., that fitting an arrow to one’s bow is a way to initiate an act of archery. Indeed, a large part of being a novice in a domain consists in acquiring knowledge of propositions about which actions are required to execute the skill and in actively considering such propositions when practicing. While Stanley and colleagues believe that experts are apt to know many *more* such propositions than nonexperts about the domain in which they possess expertise, they also accept that *merely* to increase the propositions an agent knows is insufficient to explain skill. What seems fundamentally to distinguish experts from nonexperts, on their account, is rather the practical mode of presentation under which the expert knows propositions about ways of executing the skill. It is a certain *way* of knowing propositions about which actions are required to execute a skill which experts acquire through years of disciplined practice and training. And it is presumably this way of knowing that is meant to explain why experts, but not novices, exemplify precise motor control when applying what they know in a given circumstance. On this view, to act skillfully consists in deploying one’s knowledge that w is a way to ϕ in a particular circumstance under a distinctively practical mode of presentation.

With the notion of a practical mode of presentation, Stanley et al. give the appearance of accommodating agential motor control within an intellectualist framework. However, when we examine the details of how such knowledge is meant to account for motor control, the appearance is found misleading. On examination, they do not grant that motor control is a genuinely agential (because rational) form of control, but rather deny its agential status. To show this, I’ll focus my remarks on Stanley and Krakauer (2013), which is Stanley’s most detailed discussion of how he envisions the relationship between the agent’s propositional knowledge of ways to ϕ and the fine-grained motor control the agent displays in ϕ -ing.

According to Stanley and Krakauer, skill or know how is a composite state that possesses a knowledge component and a nonknowledge component. The former is knowledge of a large number of facts about which actions are ways of executing the skill—paradigmatically, what to do to initiate the activity in a certain circumstance. They regard this component as explaining both the intelligence and control agents manifest in acting skillfully—paradigmatically, in strategic decision-making.³ The second component of skill is what they call “motor acuity”, which, following Shmuelof et al. (2012), they

³They write:

describe as “practice-related reductions in movement variability and increases in movement smoothness” (quoted in Stanley and Krakauer 2013, 8). Unlike the knowledge component, Stanley and Krakauer characterize the states and processes responsible for the motor acuity of a performance as nonintelligent and not under the subject’s direct control. They compare its agential status to digestive processes and to perceptual acuity: things, they suggest, the agent has no direct control over (Stanley and Krakauer 2013, 6). Their main aim is to show that being skilled at something requires factual knowledge about such things as, for example, the way to initiate the activity. They defend this claim by appealing to studies of amnesic subjects who show steady improvements in motor acuity using a foreign tool through daily practice with the tool. However, given their amnesia, these patients must rely on explicit instruction each day about the way the tool is used. Because the patients lack this knowledge at the start of each day, Stanley and Krakauer argue that these patients aren’t *skilled* in the activity, even though their motor acuity increases with daily practice. In their view, contemporary philosophical and empirical discussions of motor skill miss this obvious fact by mistakenly equating motor skill with its non-epistemic, non-intelligent, and non-agentive component—i.e., motor acuity.

Let’s grant Stanley and Krakauer that the factual knowledge they discuss is required for skilled actions, and let’s instead focus on the implications of their position for the motor control challenge. Of the two components of skill that Stanley and Krakauer put forward, only motor acuity seems to be directly relevant to the fine-grained control agents display when executing their intentions (Fridland 2014, 2739; Levy 2017). It is motor acuity that is claimed to explain improvements in motor variables that occur through practice and training, and Stanley and Krakauer acknowledge that merely adding to an agent’s factual knowledge—e.g., about which actions are required to initiate the activity—is insufficient to explain the performance improvements that come with skill development. But since, on their view, the motor acuity of a performance is the product of states and processes that do not directly manifest the agent’s control, this response to the motor control challenge seems to resolve into a version of the second intellectualist strategy outlined at the start. That is, it resolves into the claim that motor control does not instantiate rational control but also does not instantiate agential control. What explains the improvements in fine-grained motor control (as opposed to improvements at the level of conscious

The claim that skill involves knowing what to do and how to do it explains Aristotle’s comment about the distinctive nature of skill, which is that skills are under our voluntary control . . . If one can, in the relevant sense, start shooting arrows by fitting them into one’s bow, and one knows that one can do this, then shooting arrows by fitting them into one’s bow is under one’s voluntary control. Thus, skill possession requires a kind of knowledge, possession of which entails voluntary control over one’s actions. It is this feature of skill that explains its distinctive nature (Stanley and Krakauer 2013, 5).

strategy), on Stanley and Krakauer's view, is something that, by their own admission, is passive, nonintelligent, and brute mechanism. Stanley and Krakauer's strategy collapses into the denial that motor control processes manifest knowledge and rationality at all.

Further, if the subject's acquisition of a practical mode of presentation is meant simply to consist in improved motor acuity, then it is not easy to see how it could reflect any change or advancement in the subject's *understanding* of what is required to execute an activity. Motor acuity merely characterizes the efficiency of a passive, nonintelligent system. Consequently, one might wonder if Stanley and colleagues have genuinely described a mode of presentation. Perhaps the difficulties Stanley and colleagues face reflect a specific empirical conception of motor control that they would do better to avoid. For example, perhaps they could claim that the improvements occurring through practice and training are improvements in the agent's ability to attentionally settle fine details of movement in accordance with her intention, along the lines I've suggested. This could help them to avoid the result that the expert's practical mode of presentation on a way to ϕ is a brute causal matter, and rather a reflection of the agent's point of view in action. Might my attentional account of motor control provide a superior implementation for the notion of a practical mode of presentation?

I find this suggestion congenial. Attention could provide the practical mode of presentation under which an agent knows the way to ϕ when applying her knowledge in a particular circumstance. However, if this means acknowledging that motor control is attentionally based (in the manner that I proposed in §1.3), then Stanley and colleagues would be walking back on the claim that the agent's control is exhausted by the component of skill they link with factual knowledge and strategic decision. This is because, on my proposal, attention is distinct from the latter components of the agent's control. So, while my account might provide a way to salvage the notion of a practical mode of presentation from the difficulties Stanley et al.'s elaboration of it faces, it also seems to suggest that in applying one's knowledge in a concrete circumstance, one must actively draw on a nonconceptual form of control. That is, by allowing that attending provides the practical mode of presentation under which one knows the way to ϕ , Stanley and colleagues would seem to acknowledge that one's implementation of rational control depends on exercising of another form of control—attentional control. And this would be to abandon intellectualism about agential control.

I've discussed two versions of the claim that motor control is a species of rational control. Neither, I've suggested, presents a viable strategy for the intellectualist. With that, I propose to leave the first line of response and turn to consider the prospects of the second. This is to argue that the psychological states and processes underlying the fine-grained motor control agents display in acting falls outside the domain of agential

control.

1.4.2 Motor control as “subpersonal”

To justify the claim that motor control is nonagential, the intellectualist could argue that motor states and processes are merely a subpersonal enabling condition for the exercise of agential control rather than constitutive of agential control. According to this proposal, it is the agent’s decision to take a sip from a glass of water, for example, that marks the end of personal or agent level states and processes and the commencement of subpersonal processes responsible for overseeing fine-grained implementation details—e.g. with what velocity, force, grip size and aperture to reach for the glass. With the distinction between personal and subpersonal levels, the intellectualist can distinguish between two types of execution failure. The first type is explainable by reference to the agent’s rational perspective—e.g., to a false belief or an error in reasoning—and so genuinely attributable to the agent. The second type is a brute failure in the operation of subsystems that mediate an agent’s rational perspective and her overt behaviour. On this proposal, execution errors arising from failures of motor control as opposed to errors of belief belong to the second category. Any “gaps” that separate the agent’s decision and its motor execution can be closed at the subpersonal level without amending the intellectualist’s view of the mental states and processes that are constitutive of agential control.

The question for this proposal is: what does it mean to say that motor states and processes are subpersonal? I’ll consider two interpretations of this claim: first, that motor states and processes are unconscious; second, that motor states and processes are mere “reflex”. I’ll take each of these suggestions in turn.

1.4.2.1 Motor control as unconscious

Let’s begin with the claim that motor control is subpersonal because it is unconscious. Some think there is empirical support for this claim. To illustrate, in experiments requiring subjects to point at a target, subjects make rapid and accurate finger adjustments in response to sudden changes of target position, even when subjects report having been unaware of any change of target position or of having moved their fingers.⁴

This objection raises several important and difficult issues. I limit myself to the following observations.

⁴For an early study, see Bridgeman (1981). Milner & Goodale (1995/2006) appeal to these and other findings in support of their “two visual systems hypothesis” that dorsal stream visual processing is unconscious whereas ventral stream processing is conscious. For philosophical commentary, see Clark (2001), Campbell (2002), Kelly (2002), Matthen (2005), Briscoe (2009), Mole (2009), Brogaard (2011), Wu (2014a); Briscoe and Schwenkler (2015); and Shepherd (2016a).

First, the states underlying fine-grained motor guidance exhibit some of the central features of conscious states, even if they do not exhibit all of them. For example, these subjects adjust their fingers spontaneously in response to visually detected shifts of target position. This distinguishes the visual states underlying their finger movements from the visual states found in blindsight. In blindsight, subjects suffering from damage to primary visual cortex report no awareness of stimuli occurring within a certain region of their visual field (the “scotoma” or “blind field”), though there is behavioural evidence that information from stimuli within that region is nevertheless still being encoded in their visual system (likely bypassing the damaged striate cortex via subcortical pathways to extrastriate areas) (Kentridge et al. 1997). Although they do not spontaneously act on stimuli presented in their blind field, they can answer certain questions about the stimuli with above chance accuracy when asked to perform forced-choice tasks requiring simple discriminations. In contrast, the visual states underlying the finger adjustments of neurotypical subjects in the pointing paradigm are fluidly integrated within motor control. They thus satisfy some criteria for “access-consciousness”, including being poised for spontaneous use in service of action (in this case, in keeping one’s finger pointed at a target), though not other criteria, such as availability for introspective report and use in conscious reasoning.⁵ If one takes functional access to a state’s content as evidence for phenomenality—i.e., consciousness “there is something is to like for the subject” to enjoy (Nagel 1974)—then we have more reason to regard the visual states underlying the subject’s fine-grained finger movements as phenomenally conscious than we do in the case of, for example, blindsight.

Second, and more importantly, the above objection appears to assume that if subjects were unaware of performing the finger movement and so did not consciously intend to perform it, then that movement was not intentionally controlled by the agent. But I reject that assumption. In §1.2, I explained that one can consistently reject intellectualism about agential control and still accept *MOTIVATIONAL CONTROL*: that one qualifies as having exercised control over some property of one’s behaviour iff one’s behaviour having that

⁵For the terminology and seminal discussion of “access consciousness”, see Block (1995). See also Smithies (2011, 261) from whom I take the point that dorsal stream visual representations meet some of Block’s criteria on functional access-consciousness. However, Smithies’ takes the point in a more intellectualist direction than I wish to. Specifically, Smithies appeals to the fact that subjects are unable to offer a rationalizing explanation for why their fingers moved as evidence that while the visual states encoding target movement are functionally access-conscious, they are phenomenally unconscious (Smithies 2011, 263). An alternative would be to preserve the link between functional access-consciousness and phenomenal consciousness but reject the specific normative requirements Smithies places on phenomenal consciousness—namely, that it provide the subject with propositional justification. It could well be that the visual states devoted to guiding fine-grained movement, while phenomenally conscious, are ill-suited to provide the subject with the sort of justification Smithies emphasizes (given that their contents are, by hypothesis, unavailable for report and propositional reasoning).

property is the product of one's motivational (e.g., intentional) control. Against the intellectualist, I proposed that to qualify as having exercised intentional control over some feature of one's own behaviour, it isn't necessary that one have intended that one's behaviour possess the property. For example, it may be that although one did not intend to adjust one's finger, nevertheless the finger adjustment made in response to a perceived target shift was still controlled by one's intention to keep one's finger on the target. In that case, it qualifies as an automatic property of one's intentional performance. On my account, this makes it a property of the subject's performance that is likely the product of perceptual attention rather than of explicit decision.

My suggestion that intentional control over some property of one's behaviour does not require one to explicitly intend that one's behaviour possess that property plausibly coheres with ordinary thought about skilled intentional agency. To illustrate, consider an Olympic gymnast who skillfully executes a winning performance. Intuitively, she acted intentionally, which is why the judges were justified in awarding her the gold. She deserved the gold because it was her doing.⁶ Nevertheless, we also accept that many of the fine-grained behavioural adjustments that were critical to her success were done automatically—i.e., not as a result of deliberate strategy. There need be no inconsistency between these two claims. In particular, we can uphold both claims consistently if we distinguish the property of being intentionally guided from the property of being explicitly intended. And we can say the same about the automatic finger adjustments discussed earlier. One automatically adjusts one's finger in response to a moving target only because one has the conscious goal of keeping one's finger pointed at it.

In sum, the conscious status of the visual states and processes recruited in nonconceptual intention implementation is unclear. They exhibit some of the characteristics of paradigmatically conscious states (e.g., spontaneous uptake in service of conscious goals) but not others (e.g., availability for report). Further, we should reject the assumption that to qualify as intentionally controlled, an aspect of behaviour must have been consciously intended. Importantly, none of this helps to show that consciousness (either phenomenal or access) is irrelevant to an agent's possession and exercise of control. Indeed, in Chapter 3, I will argue that agents' *ownership* over their actions requires that they be in a certain sense conscious of what they're doing—a sense that I unpack in terms of conscious attention being deployed practically in service of goal execution. What the above discussion shows is merely that the content of the agent's guiding intention need not (and often will not) be completely determinate with respect to execution details. There are nonconceptual properties of one's behaviour that are settled nonconceptually by attending rather

⁶See also Fridland (2015) and Luthra (2016) on the relevance of practices of praising athletes for skilled performances, whose discussions are largely congenial to my own.

than conceptually through decision-making. This, I suggest, accords with everyday assumptions about skilled agency

1.4.2.2 Motor control as motivationally unguided

Let's turn to the claim that motor control is subpersonal in the sense of "mere reflex". While there is something to this suggestion, we must refine it to take account of an obvious difference between reflexes, as these are traditionally conceived, and the bodily movements that occur when we act. Specifically, whereas a classical reflex is mediated by a mandatory connection between sensory input and motor output, motor control processes deployed in action are initiated by an agent's decision to act. Despite this difference in how they are initiated, one might nevertheless insist that motor control processes are importantly reflex-like because of how they unfold. Once initiated by decision, one might suggest that skilled behaviours unfold without any further input from intention, unless perhaps halted by the agent's decision to adopt a new strategy. Instead, they follow a stereotyped routine that has been developed through practice: a tightly bundled series of responses, each step of which is triggered by a different sensory input. This would allow one to acknowledge that motor representations are responsible for making rapid online adjustments and modifications in response to incoming sensory feedback. But it would imply that these adjustments and modifications are not made in light of the relevance of the input to the agent's current goals. Instead, they are pre-set moves in a routine. In particular, on this proposal, motor control systems do not satisfy the following condition on motivational guidance:

MOTIVATIONAL GUIDANCE: A guidance mechanism—i.e., one responsible for effecting compensatory behavioural adjustments in response to changing sensory input—instantiates motivational control only if the behavioural adjustments that it makes are sensitive to the content of a motivational state of the agent.

So, while motor control systems are certainly triggered motivationally, according to this proposal, they are not guided motivationally.⁷

With this suggestion in place, the intellectualist might now draw on Frankfurt's (1978) observation that an agent's control over her own behaviour does not merely consist in the triggering of behaviour by her decision—as if, once an agent decides to perform an action,

⁷Papineau expresses a view in this vicinity when he writes:

[T]op-level batting is more like an *automatic reflex* than any consciously controlled sequence of movements. The basic facts of timing, plus the evidence rehearsed in the last three sections, all argue that the execution of a specific shot in response to the bowler's delivery is an automatic reaction honed by thousands of hours of previous practice. (Papineau 2013, 184, italics added)

behaviour unfolds wholly independently of the agent. Rather, Frankfurt points out, an agent guides a behaviour's course, standing ready to intervene to cope with unforeseen obstacles to goal attainment. To illustrate with an example favourable to intellectualism, consider our example from §1.1 of Athena navigating a foreign city in order to see City Hall. After consulting her map, planning her route, and initiating her journey, suppose Athena encounters an unexpected obstacle—e.g., a road-closure. In response, she revises and updates her original plan and decides upon a new route. By redeploying the same rational capacities that she had used in selecting and initiating action in order to deal with unexpected challenges along the way, Athena manifests a form of guidance. Specifically, she manifests *rational* guidance.⁸ Such a readiness to adapt flexibly to novel challenges in light of one's goals seems essentially to distinguish actions from behaviours that, however overtly purposive and adaptive, unfold outside the agent's control. The intellectualist might thus plausibly propose:

MOTIVATIONAL CONTROL → MOTIVATIONAL GUIDANCE: The agent exercises motivational control over some property of her behaviour only if that property of her behaviour is the product of motivational guidance.

Since, according to the intellectualist, motor control processes do not satisfy MOTIVATIONAL GUIDANCE, they do not satisfy a necessary condition for motivational control. And given MOTIVATIONAL CONTROL—i.e., that an agent exercises control over some property of her behaviour iff her behaviour having that property is a result of motivational control—it follows that they do not express agential control. Agents may be able to decide whether a motor routine as a whole will or will not occur. But if its manner of guidance is goal-insensitive, then it is not the agent who guides fine-grained movements. Rather, it is a subpersonal system within her.⁹

⁸Cf. Stanley and Krakauer:

The same kind of knowledge that is used to initiate an activity can also be injected at anytime in the ongoing course of that activity. For example, a tennis player changes her mind and switches from a groundstroke to a drop shot based on the position of the opponent (Stanley and Krakauer 2013, 5).

⁹It might be noted that while the intellectualist position we are currently considering is Frankfurt-inspired, Frankfurt himself would almost certainly oppose this implementation of the claim that the agent's control consists in a form of guidance control. For example, in reply to Bratman's claim that "deliberation and practical reasoning play fundamental roles in our agency", Frankfurt writes:

I am not so sure that deliberation and practical reasoning do play fundamental roles in our agency, if this means that they are essential to our capacity to function as agents. No doubt they figure prominently in our selection of ends and in our design of plans for reaching our goals. It is far from clear, however, that they are indispensable conditions of action. (2002, 90).

Nevertheless, it seems to me that Frankfurt's arguments for the requirements of guidance control in action are orthogonal to his objections to intellectualism.

In response, I wish to concede all of: MOTIVATIONAL CONTROL, MOTIVATIONAL GUIDANCE, and MOTIVATIONAL CONTROL → MOTIVATIONAL GUIDANCE. Motivational guidance does appear to present a plausible requirement for motivational and so agential control. If it were true that fine-grained motor processes are not motivationally guided, this would seem to show that motor control merely enables, but is not itself an exercise of, agential control.

I also want to acknowledge that some human behaviour may be routine-based in the manner described—i.e., stimulus-driven once initiated. And I even want to acknowledge that motor control processes, quite generally, are likely governed by partly autonomous principles that are less sensitive to the contents of the agent’s intentions than, for example, explicit reasoning is. The latter is suggested by empirical work demonstrating that motor control processes, especially under conditions of distorted sensory feedback, can become partly insensitive to an agent’s (full range of) goals in acting—e.g. work on visuomotor adaptation and implicit motor learning (e.g., Mazzoni and Krakauer 2006; Day et al. 2016; McDougle et al. 2016). This implies there are limits on the information motor control processes can access in their calculations, rendering them at least partially semantically insensitive to the contents of the agent’s goals and beliefs. That is, motor system calculations are, to some degree, informationally “encapsulated” (Fodor 1983; Levy 2017).

Despite the above concessions, the intellectualist’s position is too simplistic. It requires that not only some fine-grained motor control is routine-based and hence reflex-like but that *all* is. And this is too strong, as I’ll now explain.

1.4.2.3 Motor control is motivationally guided

The suggestion that fine-grained motor control is motivationally unguided runs counter to what motor control theorists have uncovered about the underlying control strategy employed in achieving coordinated movement in human subjects (Todorov and Jordan 2002; Todorov 2004).¹⁰ To see this, consider that a basic aim for a theory of motor control is to explain how, given the enormous variability in the potential ways the human body can move in order to achieve a given goal, human subjects nonetheless repeatedly and reliably manage to coordinate their movements in order to achieve the same goal across performances in highly variable environments. The question animating much work in motor control is what basic strategies humans employ in reducing the vast variability or degrees of freedom in the motor system to a precise trajectory that will lead to their goals. There are two basic answers. According to a serial model, there is an early planning phase

¹⁰Fridland (2017b; 2017a) has also drawn on the framework I’m about to discuss to argue against an oversimplified conception of “automatic” psychological and behavioural processes as exclusively stimulus-driven. My discussion here is indebted to hers. Where my discussion goes further than Fridland’s is in relating the flexible goal-directedness of motor control to attentional guidance (though Fridland also acknowledges a possible role for attention).

in which one out of vastly many movement trajectories is selected (and redundancy eliminated) and a late execution phase in which the selected trajectory is programmed and executed using relatively simple “trajectory tracking” algorithms. On these serial models, the bulk of the computational burden falls to the early decision stage, with the implementation stage being the mere playing out of inflexible routines developed over time through practice-related improvements in motor acuity. These models suggest a relatively smart, early, and flexible stage and a relatively dumb, late, and inflexible stage. Against the serial model, “optimal control theory” rejects the assumption of an early decision stage and a late implementation stage. It instead regards the agent’s task as informing fine-grained motor control directly and continuously throughout task performance.

The principal evidence in support of optimal control theory over serial models concerns the preponderance of movement variability in skilled performances like walking, grasping, etc. Serial models predict a global reduction in movement variability with the development of motor acuity through practice and training (leading to increasingly stereotyped movement patterns). By contrast, optimal control theory does not make this prediction. Instead, the optimal control strategy, according to the latter theory, is to reduce variability only where doing so is relevant to task success, leading to “task-constrained variability” in performance. As Todorov and Jordan (2002) explain, the latter prediction is empirically better supported. Even in motor tasks at which the subject is highly skilled, one finds significant variation along task-irrelevant dimensions from one performance to the next, rather than the overall suppression of variation that one would predict if a serial model were correct. Quick fine-grained adjustments are observed principally in response to environmental perturbations that affect task-relevant dimensions (e.g., obstacles to task-success). Otherwise, they remain uncompensated for, leading to variability from performance to performance. This pattern of task-constrained variability counts against the assumption that an agent’s intention figures in an early preparatory stage of trajectory selection, after which routines blindly implement the pre-selected trajectory. It instead suggests a picture on which the agent’s intention actively modulates its own fine-grained execution throughout movement execution.¹¹

My framework interprets such fine-grained, yet task-sensitive behavioural adjustments as reflecting attentional guidance. In particular, they reflect the ongoing modulatory effects of the agent’s task-set on sensorimotor systems—including, in this case, proprioceptive states that register passively caused changes in the state of the agent’s own body. On this view, movement variability between performances reflects a capacity to ignore sensory feedback not directly pertinent to current behavioural goals. Because

¹¹For complementary considerations against “serial” models of motor control primarily based in neurophysiological considerations, see Cisek (2007).

it is not worth bothering to enforce uniformity between performances with respect to execution details that aren't relevant to goal attainment, detected bodily disturbances that contribute to variation along these dimensions are attentionally de-prioritized (ignored). By contrast, if a detected disturbance in movement threatens to introduce variance with respect to a feature of one's performance where uniformity is critical for goal-success, this disturbance will be met with quick compensatory adjustments to ensure no variation in the relevant behavioural feature is in fact introduced. There is a natural role here for attending, conceived of as the direct modulation of sensorimotor systems by task-set, in realizing the kind of flexible, goal-directed motor control observed in human motor performance.

Importantly, my proposal entails that agents use attention to control their performance *throughout* intention implementation, rather than using attention only an early, preprogramming stage of motor control. If attending contributed in the latter manner, it would not be a source of motivational guidance, but of mere motivational triggering. This assumption about the ongoing contribution of attention in the guidance of motor execution receives support from eye-tracking experiments investigating the precise point at which expert golfers lose control over their swing when they “choke under pressure”—i.e., perform far below their skill level because of high levels of anxiety. Using gaze fixation duration (“quiet eye duration”) as a measure of the dwell time of visual attention on the ball (known to be longer in experts than in novices), Vine et al. (2013) compared fixation duration for successful putts and unsuccessful putts at three phases of the action: pre-backswing, during swing, and after swing. They found that the amount of time experts visually fixated the ball during the pre-backswing phase of action preparation was comparable in the condition in which they performed a successful putt and the condition in which they choked and failed. However, fixation durations in the “during” and “after” swing phases were significantly shorter when subjects choked and missed than when they putted successfully. On this basis, the authors argue that “performance failure under pressure appears to be due to disruptions in attentional control once movement has been initiated” (Vine et al. 2013, 1988). As they note, this provides empirical support to models of expertise—in particular, “attentional control theory” (Eysenck et al. 2007; Eysenck and Derakshan 2011)—according to which attention plays “an online control function, providing visual sensory information as the movement unfolds” (Vine et al. 2013, 1988). As one might have intuitively expected, expert golfers keep their eye on the ball for longer than novices do, and when they choke under pressure they are quicker to take their attention off it. This accords with my view that attention contributes a basis for agential guidance in fine-grained intention implementation.¹²

¹²In Chapter 2 (§2.5.2), I return to the example of choking under pressure and the account that attentional

I've argued that at least some states and processes underlying skilled motor control satisfy the requirements of MOTIVATIONAL GUIDANCE. Even if the states and processes governing fine-grained motor control are not always sensitive to all of the contents of all the agent's intentions in acting, they are nevertheless directly sensitive to some of these contents and are so *throughout* action. Consequently, the intellectualist's attempt to dismiss motor control as merely "subpersonal" on the grounds that this control is goal-insensitive is unsuccessful. I've further proposed, consistently with the claims of attentional control theory, that attention is the psychological basis for the guidance agents exercise in fine-grained motor control. I want now to elaborate on what I take this guidance to consist in.

1.5 Coping with distraction

To claim that agents "guide" their behaviour is to say that they oversee its progress. In doing so, they are poised to effect compensatory adjustments in response to detected disturbances that could otherwise cause their behaviour to deviate from its intended course. I've proposed in this chapter that, in the domain of skilled bodily action at least, the agent's guidance of her behaviour is realized by attention. Intuitively, agents guide their conduct by attending to what they are doing. But why exactly is attention so important for motor control? The answer I'll offer in this section is simple: because otherwise agents would likely become distracted. Having explained how I understand this claim, we'll be in a position to contrast my proposal with other recent accounts that have emphasized a central connection between attention and action.

To illustrate the problem of distraction, consider a classic paradigm from Eriksen and Eriksen (1974). In their "selective set" paradigm, subjects have the task of responding rapidly to the identity of a target character placed in the middle of a row of five characters. For example, if the character is an H or a K, subjects must respond by quickly moving a lever in one direction, but if it is a S, they are to respond by moving it the other direction. Eriksen and Eriksen found that if the subject's target is crowded by task-incongruent distractors (i.e., characters that have been paired by task instruction with

control theory gives of it. In my view, a correct understanding of the loss of attentional control that occurs when experts choke is a bit more subtle than proponents of attentional control theory typically depict it as being (though I agree with the theory's central thesis that choking under pressure constitutes a form of *distraction*). In the present context, however, the point to emphasize is not the exact character of the control loss that occurs when experts choke under pressure, but the theory's claim about the control that experts do exercise when they are *not* choking under pressure. This control is what Vine et al. (2013) (following Eysenck et al. 2007) call "goal-directed" attention. By this, they mean the kind of intention- or task-directed attention that I have been discussing in this chapter. To this extent, I agree with the conclusions Vine et al. draw from their study and, with them, take attentional control theory to provide an empirically well-supported account of expert motor control.

the opposite response to that of the target) rather than by task-congruent distractors (i.e., characters paired by task instruction with the same response as that of the target), then response times are impaired. For example, subjects are slower to respond if shown SSHSS (task-incongruent) than KKHKK (task-congruent). This is the “flanker incompatibility effect”.

The flanker incompatibility effect illustrates that distractor stimuli can present a barrier to fully successful motor control. In particular, distractors can cause “response conflicts” or “response competition”, which reduces processing efficiency and ultimately impairs performance. Another finding of Eriksen and Eriksen’s provides a clue to this problem’s solution. While task-incongruent distractors negatively impacted performance when placed within one visual angle of the agent’s target, reaction times for task-congruent and incongruent distractors were comparable when distractors were placed beyond this distance from the target. It is as if environmental stimuli falling beyond the zone of attentional enhancement are effectively suppressed and blocked from figuring in motor control processes. Intuitively, they are ignored.

Eriksen and Eriksen’s finding helped give rise to the popular notion of a “spotlight” of visual attention. In particular, Eriksen and Eriksen suggested that visual attention is organized around a retina-centric spatial frame of reference with an approximately 1 visual degree radius. Subsequent work suggests that this particular claim of theirs is probably too simplistic, and I do not want to be committed to that aspect of their proposal here. To illustrate just one difficulty, the particular spatial organization that visual attention takes in a given instance may depend contextually on the agent’s task. In the terms of my proposal here, it appears to depend on the specific kinds of perceptual and perceptuomotor states that the agent is motivationally harnessing in order to execute her goals. For example, in studies investigating distractor interference in a manual reaching task, visual attention was found not to use a retina-centric reference frame, as Eriksen and Eriksen had found, but rather a hand-based reference frame organized according to proximity to the agent’s moving hand (Tipper et al. 1992, 1998; see also Linnell et al. 2005). If visual attention is rightly understood as having a spatial frame of reference, then it is one that might vary with context depending on the specific kinds of visual states that are being attentionally enhanced. In turn, there is dispute over whether the interference caused by the nearby flankers in Eriksen and Eriksen’s paradigm reflects a spatially defined spotlight or rather “gestalt” principles of perceptual organization. For example, it may instead be that subjects perceive the crowding flankers as belonging to a single group or object as the target, leading them to be co-selected (see, e.g., Driver and Baylis 1989).

Fortunately, these questions about the organization of attention are not essential to my main point. My main point is about the broader functional role that attention appears to

be playing in response control, regardless of the specific organization attention happens to take in a given context of action. Even if one is skeptical of the suggestion that the structure of visual attention is like a spotlight, we can affirm the following points: that the perception of distractors poses a source of interference to successful action (i.e., “distractor interference”); that agents compensate for this source of interference (to the extent they are able to) by attending to what is goal-relevant and ignoring what is not; and that where agents fail to compensate for this source of interference, their control is impaired (e.g., processing efficiency declines) and they tend to perform worse as a result (e.g., are slower to respond or make more errors). Consequently, my claim that attention realizes a form of “guidance” does not merely rest with the claim that attention sustains a skilled action as opposed to merely triggering its occurrence (as discussed in the previous section). In addition, it rests on the claim that attention is a process whereby the deleterious effects of distractor interference on action are managed. And in particular I want to suggest that we can grasp these claims about the functional role of attention in the managing of distractions as reflecting a coordinative role in motor control.

To bring out the coordinative role of attention in motor control, consider the action of reaching to grasp an apple from a tree full of apples (Neumann 1987). To succeed, the same apple that informs the parameter specifications for your reach (e.g., its speed, trajectory, force, and final location) must also inform the specifications for your grasp (e.g., its aperture and force), and the same apple informing specifications at earlier phases in the reach must also do so at later phases. From an information-processing perspective, such coordination isn’t trivial. It is known, for example, that the reaching and grasping components of prehension recruit functionally distinct visuomotor processing streams.¹³ For coherent behaviour to arise, the activities of each of these component subsystems must be coordinated. Each must converge on a common environmental source when setting their respective motor parameters, rather than draw information incompatibly from distinct sources. As Neumann observes, without a way to ensure the selective “coupling” and “decoupling” of information flow among functionally separate subsystems in motor programming, action would disintegrate into “behavioural chaos”. It would become disorganized, uncoordinated, and incoherent.

Attention, Neumann proposed, is the mechanism that plays this coordinative role (see also Allport 1987). We can incorporate Neumann’s proposal into my framework as fol-

¹³As Jeannerod (1997, 21-22) explains, the reaching channel specializes in “extrinsic” properties of an object (e.g., location in space with respect to the body), functions to transport the hand to a desired location, and uses a body-centered coordinate system. By contrast, the grasping channel specializes in “intrinsic” properties like shape or size, functions to shape the hand appropriately to the object for a given purpose, and uses a hand-centered coordinate system. Further, these systems likely become engaged at distinct stages in an action (see Matthen 2005, 303 for discussion).

lows. The process of attending is realized in a biasing process that modulates lower level visual and visuomotor subsystems. The resultant state of attention is realized in the convergence of these diverse subsystems upon a common source to work together in programming an intention's execution. The suppressive or inhibitory aspect of attention serves to exclude other information arriving through the senses from contributing to motor specifications and thereby producing response conflicts. In the ideal case, the agent's intention biases her perceptual and perceptuomotor states so that only information directly relevant to a specific implementation is allowed to contribute to generating a motor representation. Here, we might think of the attention of the skilled practitioner performing her craft at her peak—e.g., the expert golfer successfully executing a putt. Because potential distractions are effectively suppressed, the golfer's movements are smooth and her adjustments to task-relevant changes in the environment (e.g., changing wind conditions) are quick and efficient.

To consolidate, let's return to the solution that I proposed to the motor control challenge in §1.3. According to that proposal, whereas agents draw on reasoning and planning in order to settle higher level questions of strategy (e.g., whether to Φ , given one's intention to Ψ), they draw on perceptual attention in settling lower level implementational questions about the precise manner in which they will move their bodies (e.g., whether to Φ in determinate manner w , given one's intention to Φ). I've argued that agents settle fine-grained, nonconceptually represented properties of their actions through the intentional harnessing of available nonconceptual information toward that end. Our discussion in this and the preceding section adds a new dimension to this picture of agential motor control. It adds that in settling how one will move, part of the control that one exercises consists in compensating for the potential interference effects caused by distracting nonconceptual information. Consequently, agents must genuinely *guide* their performances. This requires settling, on an ongoing basis, which way one will move in order to fulfill one's intention. And, in organizing one's response, one must adjust for the persistent threat to behavioural coordination that distractor interference poses. Finally, what our discussion in this section brings out is that there are clear *limits* to an agent's ability to block distracting information from interfering with response. It may be that when a distractor (e.g., a task-incongruent flanker in a selective set paradigm) is not fully suppressed, the distractor and one's target will compete and performance will decline as a consequence. In Eriksen and Eriksen's selective set paradigm, this manifests in longer response times. In other contexts, it will manifest in more execution errors. In Neumann's hypothetical attention-free scenario, it would be completely crippling. In such a scenario, we would not be able even to reach out and grasp a thing. The motor control challenge would be insurmountable (§1.2).

1.6 Comparison with extant accounts

The proposal I've outlined in this chapter about attention's role in action bears some important points of contact with other recent treatments of the topic. To further elucidate and defend my proposal, I'll conclude by contrasting my view with these extant accounts.

1.6.1 Cognitive unison theory

Mole (2011) calls attention "cognitive unison". Part of what makes unison an apt metaphor for attention is the suggestion of an absence of interference or cacophony: a coherently integrated state into which a network settles, poised to support flexible, efficient action. This accords well with my view of attending as serving a coordinative role and as the alignment of component subsystems around a common purpose. Further, we each understand the role of attention in relation to the elimination of distraction. This much I find congenial to my proposal and in step with contemporary thinking about attention.

One place where our views differ is in how they conceive of distraction. Mole develops his account of attention primarily with an eye to accommodating empirical evidence about dual-task interference. This is the processing interference that results when subjects attempt simultaneously to perform two tasks that draw on overlapping sets of processing resources, resulting in performance decrements to one or both of the tasks. His view is, approximately, that if none of the psychological resources that an agent can deploy in service of a task, given her understanding of that task, are being used to serve any *other* task, then her performance of the task will be characterized by an absence of task-irrelevant processing. Consequently, it will exemplify cognitive unison or attention.

By contrast, my account understands attention's role in relation to a more fundamental form of processing interference than dual-task interference. To illustrate, even when a subject is attentively engaged in Eriksen and Eriksen's selective set paradigm, visual processing of the task-incongruent flanking distractors (when located very close to the target) contributes to slower reaction times. In a case like this, it is possible that all the psychological resources that the subject can bring to bear on her task (given her understanding of it) either are being deployed to that task or are poised for deployment. The subject might be performing her task as attentively as anyone can. Yet, if the flanking items crowd the target, distractor interference will occur, and the subject's performance will be worse relative to the condition in which the flankers are placed further from the target and successfully ignored. I take this interference effect to reflect a genuine form of distraction. Furthermore, I take this to be a form of distraction that attention has the function of overcoming (though here attention fails fully to do so). And yet it seems to have nothing to do with distraction by another task.

To amplify the difference between Mole's and my accounts, consider the following question for Mole's account. If all of the processing resources that a subject can, with understanding, potentially apply to her current task are poised for deployment in its service, then what is to prevent all of those resources becoming active simultaneously in order to implement different possible, but mutually inconsistent ways of executing the task? I read Neumann as raising a version of this question when he asks:

The problem is how to avoid the behavioral chaos that would result from an attempt to simultaneously perform all possible actions for which sufficient causes exist, i.e., that are in agreement with current motives, for which the required skills are available and that conform to the actual stimulus situation (Neumann 1987, 347).

Since none of the processing resources that the agent brings to bear, as she plunges into behavioural chaos, are occupied with any task other than her current one (e.g., of grabbing an apple from the tree), the chaotic state Neumann describes might seem to meet Mole's criterion for cognitive unison. And yet it is a dramatic case of processing interference: only a subset of those processing resources can be deployed *coherently* at a given time. Mole acknowledges that agents typically possess multiple strategies for task implementation and that agents must select one strategy as the way they will execute the task (see Mole 2011 63 ff.). However, he attributes no privileged role to attention in the agent's coming to settle on a particular implementation strategy. My account, by contrast, proceeds on the assumption that attention does play a privileged role here—namely, in settling the specific way an agent's intention will be executed.

Possibly because Mole does not regard attending as having any necessary role to play in dealing with this more basic variety of distractor interference—a form of processing interference we might regard as emanating “bottom-up” with the sensory signal as opposed to “top-down” from the agent's goals—he is led to a relatively modest conception of attention's role in action. On his view, attention is useful but often unnecessary unless the task at hand is particularly “cognitively fragile” or “difficult” (and so “attention-demanding”) (Mole 2011, 114). Many of our skilled actions, on Mole's account, occur in the absence of any attention (Mole 2011, 62-3). On my view, by contrast, attention is deployed even in the case of relatively simple, effortless, and automatic actions like reaching out toprehend an apple.

Though I rest my account on a conception of distractor interference seemingly more basic than dual-task interference, this is not to suggest that it is a straightforward matter to determine which type of interference explains a given instance of impaired control. As we'll see in the next chapter, much of what is traditionally classified as exclusively

“bottom-up” or “stimulus-driven” processing interference—e.g., an especially “salient” distractor in a visual search task or an expert’s anxiety-induced shift of attention to the face of a threatening-looking audience member—turns out, on my view, to be directed by the agent’s motivational priorities, broadly conceived. If this is right, then a suitably generalized notion of dual-task interference (“motivational interference”) may explain more interference effects than sometimes supposed. Nevertheless, on my view, the latter type of interference is not the problem that attention exists, in the first instance, in order to solve. On the contrary, motivational interference threatens to *undo* that solution by introducing conflicts within the source states of attention.

1.6.2 The “could have done otherwise” model

Within the extant literature on attention, the account of attention nearest to my own is from Wu (2011a; 2011b; 2014b). Nevertheless, beyond genuine points of contact lie opponent conceptions of the kind of *control* agents exercise in acting. Because these differences between us are in my view quite fundamental, and because they help to motivate aspects of my investigation in later chapters, I will devote some time to assessing Wu’s account. I’ll begin by highlighting some points on which Wu and I agree.

Wu and I each hold that agents need attention, at least most fundamentally, because they can carry more information about their surroundings than they are able to put to use when acting. In this sense, we agree that attention arises originally from limitations in the capacity to act in the face of surplus information-processing capacity (rather than any intrinsic bottleneck on that capacity).¹⁴ Further, we agree in attributing to attention a role in specifying finer-grained movement parameters than are represented by the content of the agent’s intention, but that are nonetheless necessary for successful intention execution. In this sense, we both take attention to be the central source of skilled motor control.

However, Wu and I situate these points within opposing conceptions of agency. In particular, Wu understands agency fundamentally in relation to what he calls the “many-many problem”. In his terms, agents face many perceptual “inputs” and many potential behavioural “outputs”. This yields a many-many mapping between available inputs and potential outputs which Wu calls the agent’s “behavioural space”. For action to occur, Wu claims, the agent must select one of these many input-output mappings—i.e., a “path” through behavioural space. Attention, Wu suggests, is the agent’s selection of a behavioural path. Because every action requires a solution to the many-many problem and only attention can solve this problem, Wu concludes that every action requires

¹⁴Cf. also Allport (1987) and Neumann (1987). Note that none of these authors deny that agents possess information-processing limitations. They instead deny that attention’s most basic role is to manage such limitations.

attention. Attention is “selection for action”.

What justification does Wu provide for the claim that every action poses a many-many problem? The answer is to be found in the background conception of agency on which Wu builds his account:

... [F]or agency to be possible, there must be behavioral options, even if it is just the option of not acting. A behavior space must open up with more than one path, and now there is a Many-Many Problem ... All actions then emerge from an appropriate behavior space where actions entail ... choosing one among other behavioral paths. This “choice” is not available in pure reflex. (2014b, 90)

In other work, Wu motivates this view by inviting us to consider a possible world in which creatures do not confront many-many problems. We will be forced to conclude, he suggests, that agency is absent from that world. Regardless how overtly complex and coordinated the movements of creatures in that world might be, these creatures are not *agents* because “the presentation of possibilities is denied them” (2011b, 54). For Wu, it is a creature’s selection of a path through a space of alternate behavioural possibilities that constitutes “the critical moment in agency” (2014b, 81). And it is that selection of one from many alternate possibilities (even if just the possibility of refraining from action) that, on Wu’s account, essentially distinguishes the actions *we do* from the “pure reflexes” that we passively undergo. When the doctor taps your knee and your leg jolts, you may say “I didn’t do that!” What makes your assertion correct (if it is), Wu claims, is that you could not have done otherwise than to move your leg. So, moving it wasn’t up to you.

This all suggests that Wu endorses the following:

REGULATIVE CONTROL: Agency consists in the exercise of the power to select one out of many alternative behavioural possibilities (minimally, the power to Φ or not to Φ).

In more detail, an agent exercises “regulative control” over an action iff the agent selects it from a set of alternative possible responses remaining open until the time at which she acts. The agent thereby “regulates” between opposing behavioural alternatives.¹⁵ If this explains the basis of the agent’s control in action, then without the freedom to select and act otherwise than one does, the agent’s conduct wouldn’t be *up to* the agent. The fundamental problem of agency, on this conception, is to select which alternative behavioural

¹⁵I borrow the terminology of “regulative control” from Fischer (1994). Fischer opposes regulative control with “guidance control”, the latter of which does not require alternative behavioural possibilities, but concerns the actual causal mechanism responsible for behaviour. Fischer is, in turn, elaborating distinctions in Frankfurt (1971).

possibilities to actualize given all the information the agent has available and her full repertoire of behavioural capacities. If attention is the way agents select between alternate possibilities, attention is revealed as the psychological basis of agents' exercises of regulative control. Whereas the agent selects which action to perform and on which object, the agent doesn't select a reflex because no many-many problem exists for reflexes. If reflexes are "selected" at all, they are selected subpersonally and so are outside the agent's control. Given that the distinction between action and reflex provides a mutually exclusive and jointly exhaustive classification of behaviour, attention is revealed as necessary for action of any sort.

We should resist Wu's view of attention as the exercise of regulative control in action. To begin with, the account seems to face a dilemma concerning the agential status of selection for action. In particular, either selection for action is an action or it is a reflex. If it is an action, then we face a vicious regress. This is because if selection for action is itself an action and every action requires selection for action, then selection for action will itself require selection for action. If selection for selection for action is itself an action, then it too requires selection for action, and so on *ad infinitum*. Selection for action cannot be both an action and a precondition for action of any sort.

If, on the other hand, selection for action is a reflex, then there is no role for the agent in generating behaviour. In general, if ϕ -ing is a reflex, then, Wu supposes (plausibly), it isn't the agent who ϕ -s. For example, if an arm rising is a reflex, then it is not the agent raising her arm. Similarly, if selection is a reflex, then it is not the agent selecting the response. But action requires that it be the *agent* selecting response. So, if selection is a reflex, then the selected response isn't an action.

So, whether selection for action is an action or a reflex, action is revealed to be impossible.

There are ways to avoid the above dilemma. Arguably, the most promising way to do so (and Wu's own preferred response) is to claim that the agent's selection of a specific behavioural path is not any event *in addition to* the agent's action—whether an additional action or a passive reflex. Rather, on this view, to select *just is* to act.¹⁶ On this elaboration, we do not understand the agent's selection as an event that causes one path through

¹⁶Wu expresses sympathy for this conception of the agent's selection in Wu (2014b, 97-99) and comes out more forcefully in support of it in Wu (2016):

'Selection' might suggest to some readers that the agent must do something else (selecting) in order to act. In fact, the idea is simpler. We have appropriate selection when a subject's perception of the environment is coupled to and thereby informs the production of a response. That is, selection is just a necessary product of taking a specific path in behavior space. Where we have a specific aspect of the subject's experience operating in this way, we have selection of what is experienced for action. This just is, I claim, a form of attention (2016, 108).

behavioural space to occur; rather, we understand the agent's selection as the agent's *taking* a specific path through behavioural space where alternative behavioural possibilities existed. In skilled bodily action, for example, we can think of the agent's "selection" as consisting in a certain perceptual state being used to inform a certain motor response, where alternate perceptual-motor couplings were possible. The agent's selection is not a further event leading up to this state of affairs, but the resultant state of affairs.

This does seem to provide a coherent response to the dilemma that I have posed. However, it is worth pausing to reflect on its implications. The most natural way to interpret the above response is to construe the agent's "selection" as entailing the exercise of a causal power by an agent—specifically, the power to cause one behavioural option (i.e., input-output mapping) to be actualized rather than another one. On this reading, what distinguishes the input-output coupling that occurs in action from the input-output coupling that occurs in pure reflex is that, in action but not reflex, the input-output coupling is the *agent's doing*. It is caused by the agent. Were an input and output to become coupled without it being the agent who causes them to couple—i.e., who selects the coupling—then this would just be pure reflex.

If this is right, then the proposed solution to the dilemma requires regarding the *agent* as one of the relata in an action-constituting causal relation. Non-reductive agent-causalists will be happy with such a picture.¹⁷ However, Wu is not such a theorist. In keeping with the aspirations of "reductive", "event-causal" theories of action, Wu instead means to *reduce* the agent's causal role in action to that of certain mental states, events, or processes. Specifically, his proposal is to identify the agent's role in action with the selection of a behavioural path. And this reductive aim sits uncomfortably with the dilemma-averting implementation of the selection for action account. Confronted with the question of what distinguishes the input-output coupling that occurs in action from the one occurring in pure reflex, Wu cannot appeal to the causal role of the agent. Nor does he wish to appeal to the agent's intentions or other motives as constituting the agent. This is because he maintains that it is possible for agents to attend and act in the absence of

¹⁷Steward's (2012) account of action provides an instructive point of comparison here. Much like Wu, Steward holds that, with time's passage, the agent continuously resolves the fine details of her bodily movement in a way that is highly constrained, but never fully determined, by her motivational states (e.g., her intentions). It is in the settling of such fine details that Steward finds a critical indeterministic role for the agent in action—a role the agent would be deprived of if her behaviour were determined (motivationally or otherwise) down to its fine details. Without the power to settle the answers to open questions about how one's body will move, Steward believes that agency would be abolished (much like in Wu's hypothetical world of creatures who do not confront many-many problems). Steward couches these ideas within a framework of primitivist agent-causalism. As I read her proposal, it not only requires that we reject causal determinism (which Steward is candid about), but arguably also physicalism, given its reliance on strong metaphysical emergence about the agent and the agent's causal powers in relation to her body. Assessment of Steward's rival implementation lies beyond the scope of my investigation. I highlight it mainly to note that the stakes are high here.

any motivation to do so, which he takes the case of purely “bottom-up” attention capture to demonstrate (Wu 2011a, 101; 2014b, 91-93). His view is that our motivational states may (usually) constrain how we act, but they are not *necessary* for action (cf. Steward 2012, 34, 66). In order to uphold his proposed solution to the dilemma, then, it seems Wu must retreat to the bare counterfactual claim that whatever input-output mapping in fact occurs when an agent acts *could have been otherwise*, whereas the input-output mapping that occurs in pure reflex is (in some unspecified sense) necessary or mandatory.

However, this is an unattractive way to distinguish action and reflex. For one, the proposal faces the difficult question of distinguishing the kind of indeterminism that characterizes the actions of agents from other potentially indeterministic events that are presumably not the actions of any agent, such as the specific path that water happens to take while streaming down a hill or the precise moment at which a radioactive atom emits a particle. Without invoking the agent or some event of decision as the cause of the indeterministic event, there appears to be little to distinguish the two scenarios. Consequently, Wu seems to face a particularly acute version of the “Disappearing Agent Problem” (Velleman 1992; Pereboom 2014).¹⁸

I’ll now try to pinpoint where, in my view, Wu goes astray. We each agree in emphasizing perceptual attention as the psychological capacity underlying an important form of flexibility manifest in skilled action. In particular, we agree in regarding perceptual attention as the non-deliberative means whereby agents resolve fine-grained properties of behaviour. However, by situating these ideas within the framework of REGULATIVE CONTROL, Wu provides a broadly libertarian gloss on the flexible control manifest in such actions. What most fundamentally distinguishes skilled action from pure reflex, Wu thinks, is a kind of response-freedom: the fact that sensory inputs do not “coerce” behaviour. To accommodate this, he postulates some event, process, or state of personal level selection to regulate between alternate options, and he regards occurrences of such selection as constituting a behavioural episode as an action. It is this interpretation on the flexibility exemplified in skilled action that lands him in the above difficulties. If it is the agent who

¹⁸Admittedly, Wu has in some work emphasized a type of agential control linked with intention, acknowledging that a purely automatic or unintended action would occur without agential control and so would be “passive” (e.g., 2013; 2017). However, even in these contexts Wu maintains (a) that the many-many problem provides the criteria for distinguishing action from reflex, and (b) that the many-many problem can be posed and solved (through attention) outside any context of intention. These claims entail the possibility of passive actions. In response, Wu is here backtracking on the intuitions upon which he drew to distinguish action and reflex in the first place: only the former, recall, are “up to me” or “my doing”. If he attempts to capture “up-to-me-ness” with attention and the many-many problem, then it is unclear in what sense the action is passive or beyond the agent’s control when unintended. Second, my argument in the text remains as a challenge to the claim that a genuine form of action has been identified with the many-many problem. And third, there are no passive actions in my view. While there can be impulsive, emotional, and habitual actions, every action is the agent’s doing.

selects, as Wu claims, then either this selection must be an action the agent performs or it must somehow constitute the agent as acting. The former leads to a regress, while the latter lacks an informative account of what makes the selection (e.g., a certain input-output mapping) agential. Wu's view consequently cannot provide an adequate account of the role of the agent in action.

By contrast, I endorse **MOTIVATIONAL CONTROL** as an account of the control agents exercise in acting (§1.2), which I've further specified as **MOTIVATIONAL GUIDANCE**. As the point is often put in the free will literature: whereas regulative control concerns the agent's alternate behavioural possibilities at the time of action, guidance control concerns the actual causal mechanism underlying the agent's behaviour (Frankfurt 1971; Fischer 1994). Actions are motivationally guided processes. Rather than response freedom, therefore, I attribute a necessary role to the agent's *goals* (or as I will call them in Chapter 2, the agent's "priorities") in distinguishing action from reflex. Given that attending is itself a process that unfolds under the direction of the agent's goals, attending qualifies as an action as well. It is the goal-directed process of selecting and structuring one's perceptual and perceptuomotor states to be usable in goal execution. This process need not (and arguably often will not) be itself attention-guided. Commonly, we do not attend to our attending, though the capacity to deliberately guide (and so decide) how we will attend could become important under conditions of motivational conflict. But that would be a more sophisticated form of cognitive control than we have been examining in this chapter. It is probably more common that our attention is immediately and automatically guided by our goal-setting motivational states, rather than mediated by separate acts of attention. Since I deny that acts of attention must themselves be attention-guided, no regress arises on my view. The precondition for action is not guidance by attention, but guidance by a motivational state. One important form such guidance can take is attentional, but it is not necessary that guidance take this form.

On my view, the informational challenge confronting agents is not to select a path through behavioural space and to solve many-many problems. Rather, the informational challenge confronting skilled agents is to implement coordinated and coherent behaviour in pursuit of their goals given the threat of distraction. These are different problems. In principle, an agent's motivational states could fail to constrain which sources of information the agent draws on in response selection, with the result that her behaviour becomes chaotic, disorganized, and incoherent. So long as the agent's responses are formed on the basis of a subset of the information available to her, she would qualify as traversing behavioural space, and so, on Wu's view, as exercising agential control. However, her behaviour will have been motivationally unguided. Specifically, her behaviour will have

been distracted, and so, on my view, would not exemplify the agent's control.¹⁹

For these reasons, the move to a guidance framework of agential control seems to me an improvement over the rival framework Wu chooses to work with. However, a version of the same challenge that Wu faces might be thought to arise also for accounts that purport to explain action in terms of guidance of behaviour by a motivational state of the agent. In particular, one might ask what could justify reductively identifying the *agent's* guidance of behaviour with guidance by a *motivational state*? The suspicion that nothing *could* justify doing so seems to generate a new version of the Disappearing Agent Problem, this time for my account. I address this challenge in Chapter 3. There, I suggest that attention (specifically, conscious attention) plays an essential role in answering this challenge on behalf of reductive-causal theories of action. Rather than attempting to reduce the agent's role in action to attention, as Wu proposes to do, I argue that to solve the Disappearing Agent Problem we must identify the agent with the psychological state or states serving to *direct* conscious attention.

Before we see how I propose to answer the Disappearing Agent Problem, however, we must first address a different challenge to my account. I have proposed that attending is closely tied to acting: attention *just is* the modulation of information processing by motivational states. But, the objection runs, surely there are cases where we attend without acting—cases where attention is “captured” from without rather than driven from within. The next chapter develops my proposal in a way that addresses this concern.

¹⁹This point suggests that regulative control is not sufficient for agential control. “Frankfurt cases” can be adapted to show that it is not necessary (even over very fine-grained features of one's behaviour). See Steward (2012, Ch.7) for discussion of many permutations on Frankfurt cases (though I depart from Steward's negative verdict about the efficacy of these cases).

Chapter 2

Actionism and Attention Capture

2.1 Introduction

Chapter 1 identified one way exercises of agency depend on attention. The suggestion was that finer-grained properties of action are the nonconceptually represented outputs of perceptual attention. This is in contrast to higher level properties of an action that are represented by an intention, usually as a result of practical reasoning. Given agents exercise motor control in skilled bodily action and attention is the process whereby agents exercise this control, attention is required for skilled bodily action.

In this chapter I turn to ask whether attending is *always* active. According to what I'll call "actionism" about attention, to attend is necessarily to exercise an agential capacity. This may be because attention is itself essentially an action or activity (Watzl 2011; 2017; Jennings 2012; Carruthers 2015), because attention is essentially a manner of performing an action (Mole 2011; Koralus 2014), or because attention is essentially an aspect of acting (Wu 2014b). Actionists do not merely claim that every subject of attention is also an agent capable of action. They commit to the stronger claim that *whenever* one is attending one is also acting.

Now, if attention ever occurs without action, the relationship between attention and action is contingent rather than necessary, and actionism is false. So, to refute actionism, it's sufficient to identify a single instance of attention occurring without any action being performed. Common sense suggests that there are such instances. For example, while working quietly at your desk, you hear a fly buzzing around your head. The buzzing sound *draws* or *grabs* your attention, distracting you from your work. On the face of it, automatically reorienting attention to the fly is neither an action nor serves any action that you're attempting to perform. Instead it seems like a reflex—analogueous to the way your knee jerks when tapped. If so, then there are counterexamples to actionism. This is the *objection from attention capture* to actionism.

Actionists are aware of this challenge (Mole 2011, 53; Jennings 2012, 541; Wu 2014b, 90-92; Carruthers 2015, 145 ff). However, extant responses have proceeded as if there is only one constraint on an adequate solution when really there are two. The first, more widely recognized constraint is to show how, even in cases we would intuitively describe as “captured” attention, the agent nonetheless still qualifies as active. The second, hitherto unrecognized constraint is to explain why, in some cases of capture, attention fails to fulfill its proper functional role and consequently is defective as attention. In this chapter, I defend an actionist reply to the objection from attention capture that satisfies both constraints. I will now say more about the motivations behind each of them.

The reason for the first constraint is clear enough. If one holds that one acts whenever one attends, then one must show that, even in intuitively passive cases like attention capture, the agent is active. Actionists have made various suggestions here. For example, it has been suggested that in these cases the agent is still selecting one out of many alternate behavioural possibilities (Wu 2011a, 101; 2014b, 91-3); maintaining veto power (Jennings 2012, 541); using her understanding to accomplish a task (Mole 2011, 52-3); responding to an unconscious decision (Carruthers 2015, 145 ff) or a primitive urge (Watzl 2017, Ch. 6) to attend to the attention-grabbing stimulus. In different ways, each proposal hopes to show that attention, even at its most intuitively passive and involuntary, retains an active element.

The reason for the second constraint may seem less obvious. Whether one is an actionist or not, one might suggest that what succeeds at capturing our attention often *merits* our attention. Attention capture is functionally beneficial because significant events are registered and attended to regardless of what we happen to be focusing on at the moment of capture. For example, while working at your laptop, suppose you notice a rodent-shaped entity scurrying by the floor board in the visual periphery. You halt what you’re doing to investigate what you saw. The attention you were paying to your laptop has been interrupted and automatically reallocated to where you detected motion. One might plausibly argue that here attention operates as it should: sometimes attending properly requires dropping one activity to focus on something of higher behavioural priority or importance to the agent.

This is true. But not all cases of capture are like this and so the constraint remains. Recall the example of the buzzing fly. Suppose that you have no interest in the fly and are quite certain that it poses no danger to you. It is irrelevant to you, yet it grabs your attention. In this case, your perception of the attention-capturing stimulus need not prompt you to switch from one task or project to another, more important one. Rather, the fly plainly *interferes* with your goals. It distracts you. Let’s call this a case of “distracted capture” to distinguish it from a case in which attention capture coincides with the agent

switching to a higher priority task. If, finally, the proper functional role of attention is to prevent interference from distraction (see Chapter 1, §1.5), cases of distracted capture seem to be cases of defective attention.

We can summarize the argument for my second constraint as follows:

1. The proper functional role of attention is to prevent the agent from becoming distracted by goal-irrelevant information.
2. In cases of distracted capture, the agent becomes distracted by goal-irrelevant information.

Therefore,

3. In cases of distracted capture, attention's proper functional role is not satisfied.

If this argument is successful, then there is an analogy between distracted attention and false belief. Supposing a belief's role or aim is to track the truth (see e.g., Williams 1973; Velleman 2000a), then a false belief is unsuccessful insofar as it doesn't achieve that aim. Analogously, according to the above argument, insofar as cases of distracted capture fail to achieve the aim of attention, they are unsuccessful cases of attention. And supposing, with the actionist, that attention is the exercise of an agential capacity, distracted capture is a defective exercise of that capacity. Extant actionist responses to the objection from attention capture fail to recognize this point when they depict all cases of attention as if they were equally good according to the standard of success set by attention's proper functional role. This is an error. It distorts our understanding not only of attention but also of a central range of cases in which agents lose control over their conduct. As we'll see, defective attention begets defective action.

The chapter proceeds as follows. In §2.2, I introduce the distinction between automatic and volitional deployments of attention, as this is operationalized within the empirical literature on visual attention. In §2.3, I discuss a traditional view of automatic attention as governed by *perception of environmental salience*—the “environmental salience view”—and I explain why the view is inconsistent with actionism. In §2.4, I argue against the environmental salience view. In §2.5, I motivate the rival *priority view* of automatic attention and explain how the priority view allows us to satisfy both constraints on an adequate response to the objection from attention capture. I conclude by contrasting my account with two extant actionist views of attention, each of which fails to acknowledge the possibility of defective attention.

2.2 The distinction between voluntary and automatic attention

Common sense has it that sometimes we voluntarily control how we attend. Thomas Reid writes:

Every man knows that he can turn his attention to this subject or to that, for a longer or shorter time, and with more or less intensesness, as he pleases. It is a voluntary act, and depends upon his will. (Reid 1872/2000, 127)

It is equally part of common sense, however, that attention is not always a voluntary act. As Reid explains:

It is well known, that things new and uncommon, things grand, and things that are beautiful, *draw* our attention, not in proportion to the interest we have, or think we have in them, but in a much greater proportion. Whatever moves our passions or affections draws our attention, very often, more than we wish. (Reid 1872/2000, 128, italics added)

Similarly, William James distinguishes a variety of attention requiring effort to initiate and sustain and implicating the agent's will from a form of attention that operates passively. As he describes the latter:

In *passive immediate sensorial attention* the stimulus is a sense-impression, either very intense, voluminous, or sudden . . . or it is an instinctive stimulus, a perception which, by reason of its nature rather than its mere force, appeals to some one of our normal congenital impulses and has a directly exciting quality ... [T]hese stimuli differ from one animal to another, and what most of them are in man: strange things, moving things, wild animals, bright things, pretty things, metallic things, blows, blood, etc. (James 1890/1950, 416-7, italics original)

Reid's and James' remarks about active and passive attention are intuitive. However, there are places where they do not say quite as much as we might like them to. For example, when Reid writes that things draw our attention "not in proportion to the interest we have, or think we have in them, but in a much greater proportion", Reid remains non-committal about whether it is actually the agent's interests that determine what draws attention. And when Reid claims that "whatever moves our passions or affections draws our attention", he leaves open whether it is our passions or affections that cause attention to shift or whether our passions and our attention are moved together by a common cause. I take these to be questions about the role the agent's motivational states, broadly construed—including the agent's interests, impulses, passions, and affections—play in

intuitively involuntary and automatic shifts of attention. Similar remarks apply to James, who does not clarify the role of the animal's congenital impulses in determining the focus of attention. Reid's and James' hesitancy to pronounce on these issues may be evidence that common sense and introspection alone do not resolve them. Further, because Reid's and James' distinction between two types of attention is ultimately an empirical matter, we might ask how well they stand up to empirical investigation. To verify if their claims are correct, we must consult the relevant empirical evidence.

Fortunately, there is a large body of empirical work bearing directly on Reid's and James' claims. I will show that this work ultimately vindicates Reid's and James' distinction between two kinds of attention and (more controversially) enriches that pre-theoretic distinction in ways that help to resolve whether attending is the exercise of an agential capacity. Because the vindication claim is less controversial than the enrichment claim, it makes sense to begin by examining the case for the former. Once we have seen how, in broad outline, empirical psychology supports Reid's and James' intuitive picture, we will be in a better position to consider more nuanced questions about the possible motivational dimension of automatic attention. For the remainder of this section, I'll explain some of the main paradigms psychologists have used to empirically investigate the control of attention and how these have been used to establish a distinction between two modes of attentional control. This will provide the background against which more controversial questions can be posed.

An influential experimental paradigm in the study of attentional control is the "spatial cuing paradigm" (Posner 1980). In this paradigm, the subject visually fixates at a central location on a visual display (e.g., a cross). The subject's task is, without moving her eyes from the center, to respond as quickly as possible to the appearance of a target on the left or right hand side of the screen. Before the target appears, the subject sees a cue for 100 milliseconds. Depending on the trial, the cue will either accurately predict target location (i.e., will be "valid"), inaccurately predict target location (i.e., will be "invalid"), or will not carry any information about target location (i.e., will be "neutral"). Shortly after the cue disappears, the target will appear and the subject must quickly report that she saw it or report some feature of the target. Posner (1980) found that valid cues facilitate, while invalid cues impair, target detection. Subjects are faster to detect the target when they receive a valid cue than if they receive a neutral cue, and they are slower when they receive an invalid cue than if they receive a neutral cue. This is taken as evidence that visual attention shifts (covertly) to the cued location.

Posner's spatial cuing paradigm provides a way to investigate the dynamics of visual spatial attention. It has been adapted to investigate the dynamics of auditory attention (e.g., Spence and Driver 1994) and cross-modal attention (e.g., Spence and Santangelo

2009). For example, an auditory cue at a certain location has been shown to facilitate response to a visual target at the same location and vice versa. I will focus on the visuo-spatial case.

Significantly, researchers have found that attention behaves differently depending on the type of cue the subject is shown (Jonides 1981; Posner and Cohen 1984). An “indirect” (or “central” or “symbolic”) cue is an arrow (or some other meaningful symbol) that appears at the center of the display where the subject’s gaze is fixed and symbolically indicates which side of the display the upcoming target is likely to appear. Upon seeing a central cue, subjects strategically shift their attention to the symbolically indicated location to prepare for the target’s arrival. By contrast, a “direct” (or “peripheral”) cue is a stimulus, such as a brief flash or other abrupt onset stimulus, that appears on either the right or left side of the display while the subject’s gaze is fixed at the center. In valid conditions, the subject’s target later appears where the cue had previously appeared, while in invalid conditions the target appears on the opposite side of the display to where the cue had appeared. Even when subjects know that the direct cue is a task-irrelevant distractor and explicitly form the intention to ignore it, valid peripheral cues improve performance and invalid cues impair it. This has been taken to show that a peripheral cue, unlike a central one, directs the subject’s attention “involuntarily” and “automatically”—i.e., without the subject intending to redirect her attention to the cued location. In these circumstances, researchers typically speak of the subject’s attention as having been “captured” by the cue. They also speak of subject as “distracted” by the cue (also sometimes called a “distractor”). Because one form of attention is controlled voluntarily and strategically, while the other form occurs involuntarily and automatically, the distinction is sometimes drawn between “controlled” or “voluntary” attention and “automatic” or “involuntary” attention.

These two types of attention are known to differ in their temporal characteristics. First, they differ in how much time they require to reach a cued location. Whereas a valid indirect cue maximally facilitates task performance with a cue-target onset asynchrony of 300 milliseconds, a valid direct cue does so with a cue-target onset asynchrony of about 100-120 milliseconds (Carrasco 2011). Involuntary attention thus reaches a cued location faster than voluntary attention, possibly reflecting the greater amount of time required for central cognition to become active compared with automatic processes. Second, whereas subjects can voluntarily sustain their attention to an indirectly cued location for as long as they choose, attention to a directly cued location decays rapidly (Posner and Cohen 1984). This can be demonstrated by varying cue-target onset asynchrony for direct cues. If the target appears approximately 100 milliseconds after the direct cue disappears, the cued location receives maximal facilitation relative to uncued locations. If, however, the

target appears 200 milliseconds or more after the direct cue disappears, the reverse is found to occur: subjects are now slower to detect the target at cued than at uncued locations. After an initial period of attentional facilitation, attention is immediately inhibited or suppressed at a directly cued location. This is called “inhibition of return”. No such attentional inhibition is found using an indirect cue: subjects can voluntarily maintain attention to the cued location long after the cue has disappeared. Because of these differences in temporal profile, psychologists sometimes refer to the type of attention engaged by a direct cue as “transient” attention, and the type of attention engaged by an indirect cue “sustained” attention.

In sum, common sense is right to distinguish two types of attentional deployment. Attention can be voluntarily deployed in light of one’s current task or deployed involuntarily and automatically. Though this helps to vindicate part of our intuitive thought about attention, it bears emphasis that nothing that we’ve seen so far directly resolves whether automatic attention is a counterexample to the thesis of actionism. We’ve said no more than Reid and James already had about what role, if any, the agent’s motives might play in explaining why attention behaves as it does when automatically deployed. Even if automatic attention can observably conflict with the agent’s explicit intention, it remains to be seen that it is controlled independently of the agent’s motives. To make progress with this latter issue, we must go beyond the bare distinction between voluntary and involuntary attention and examine specific proposals about the mechanisms governing how attention is allocated. Since it is the automatic or involuntary control of attention that bears most directly on the objection from attention capture, I’ll henceforth be concerned with the mechanisms of automatic (i.e., involuntary, transient) attention. The first proposal that I’ll consider explains deployments of automatic attention in terms of a stimulus property called “salience”. Properly understood, I’ll suggest, the salience view of automatic attention is inconsistent with actionism in any form.

2.3 The environmental salience view

Why do subjects attend automatically to the location of a direct cue upon seeing one in Posner’s spatial cuing paradigm? A traditional answer to this question appeals to the cue’s visual *salience*. However, if “salience” is to do the required explanatory work, it cannot simply mean the property of being such as to cause attention to shift automatically to the cued location when the cue is perceived. Understood in *that* way, an answer in terms of cue salience does not tell us anything we did not already know. It would not tell us what it is about the cue, or about ourselves, that causes attention to shift. However, there is another sense of “salience” that promises to provide a more informative answer

to our question. And this brings us to the account of automatic attention that I will be considering in this section and the next one: the environmental salience view.

According to the “environmental salience view”, automatic attention is controlled by the perception of *environmental salience*. The view has two central assumptions:

- (i) Salience is a property of the subject’s environment to which subjects are perceptually (e.g., visually) sensitive.
- (ii) In normal perceivers, the perception of salience at a location is able to cause attention to shift automatically to the location where salience was (and perhaps still is) perceived.

Assumptions (i)–(ii) are at the very least implicit in the terminology psychologists frequently use in discussing automatic attention. For example, paradigmatically attention-driving stimuli—e.g., abrupt onsets in a spatial cuing paradigm and feature singletons in a visual search paradigm—are commonly called “physically”, “objectively”, and “intrinsically” salient, and the ensuing attention-shift is called “exogenous”, “stimulus-driven”, and “bottom-up”. This suggests that salience is understood as an objective, physical feature of stimuli, perception of which is sometimes sufficient to drive attention to the salient location.

Assumptions (i)–(ii) are also explicit in traditional “salience-based” accounts of automatic attention. According to one influential model—first introduced by Koch and Ullman (1985) as a hypothetical mechanism to explain the dynamics of attention when searching for a target in a crowded visual display and later modelled computationally by Itti and Koch (2000; 2001)—there is a “salience map” in the human brain governing how attention automatically evolves over time. This is hypothesized to be a topographically organized map that integrates multiple “feature maps” into a single representation of the relative conspicuousness (or salience) of each location in a scene. Salience is computed using image-based algorithms that measure how dissimilar each location is from neighbouring locations, synchronically and diachronically, along multiple dimensions of variation (e.g., colour, orientation, motion), then aggregate these individual measurements into a single value representing overall salience. The location with the highest salience value then drives attention in a “winner-takes-all” fashion and becomes a likely target for eye movement. Once attended, the salient location is then promptly suppressed (inhibition of return). This is hypothesized to endow attention with an “internal dynamics” so that it can scan a scene for novel information without getting stuck cycling between only the most physically salient locations.

Consistent with assumption (i), salience-based models regard salience as a measurable property of stimuli to which the subject is visually sensitive in virtue of a salience map.

Consistent with assumption (ii), perceiving environmental salience can cause an automatic attention shift. For example, a red square in a visual display full of green squares would be considered especially salient because its location differs from its surroundings more than any other location in the display—here, with respect to the dimension of hue. Consequently, the red item will visually “pop out” and automatically summon attention to its location, resulting in very fast response times in a visual search task (Treisman and Gelade 1980). This is an example of “static” discontinuity. Similarly, Yantis and Jonides (1984) and Jonides and Yantis (1988) proposed to explain the tendency of attention to shift automatically to the location of a direct or peripheral cue in Posner’s spatial cuing paradigm by the fact that the cue in these experiments has abrupt visual onset (e.g., a sudden flash). This is an example of “dynamic” discontinuity because the dissimilarity is between a location at one moment (just before the cue arrives) and the same location at a subsequent moment (when the cue arrives). Because abrupt onset involves dissimilarity along multiple featural dimensions, it was suggested to be especially salient and able to redirect attention independently of behavioural context. According to the environmental salience view, visual attention unfolds automatically by continually being captured by perceptions of environmental salience.

Returning to the thesis of actionism about attention, I want to suggest that the environmental salience view of automatic attention should make actionists uneasy. Actionism, recall, is the view that every episode of attention manifests an agential capacity, either because attention is, by its nature, an action, a manner of action performance, or a component of action. By contrast, the environmental salience view seems to make automatic attention more similar to a passive reflex in which your body moves but it isn’t you who moves it. When environmental salience captures your attention, it seems it isn’t you who moves your attention to a new location, but the salience of the location. Like a knee-jerk reflex, salience-driven attentional reorienting seems to be fundamentally passive.

To consolidate this diagnosis of conflict between actionism and the environmental salience view, consider a sample of recent actionist proposals for how to accommodate the intuitive notion of captured attention:

- When something salient captures attention, the agent selects the salient stimulus for some action (e.g., orienting). The agent is active in selecting because it was psychologically possible for her to select a different stimulus for response or no stimulus at all (Wu 2014b, 90, 92).
- When something salient captures attention, the agent begins to attentively perform a different task from the one she was performing when she first perceived the salient stimulus. The agent is active because her task-performance is guided by her under-

standing of the task (Mole 2011, 52-6).

- When something salient captures attention, the agent's act of attention is governed by a motivational state, such as a decision (Carruthers 2015) or a more primitive urge-like state (Watzl 2017, Ch. 6) to turn her attention to the salient stimulus. The agent is active in shifting attention because the attention-shift is controlled by a personal level motivational state.

Common to each of these proposals is the suggestion that attention, even when captured automatically, is controlled by something internal to the agent: a selection among alternatives, an understanding of how to achieve a certain task, or a motivational state like a decision or an urge to attend to the salient item. By contrast, the environmental salience view leaves no room for such internal variables because it regards automatic attention as governed purely by the perceived qualities of the stimulus. Put differently, each actionist proposal takes attention, including capture, to be controlled at least partly *endogenously* whereas the environmental salience view takes attention capture to be purely *exogenous* or *stimulus-driven*. It is no coincidence that each of the actionist proposals sketched here disagrees with the latter claim. The latter implies that attention, when captured, is controlled by something external to the agent, which is plausibly the antithesis of agential control. While I've considered only a sample of actionist proposals, others are apt to be inconsistent with this feature of the environmental salience view too.

It is worth noting that all parties will agree that attention capture is "automatic" in that it occurs without being intended. On its own, this is consistent with actionism because, as argued in Chapter 1, many automatic processes manifest agential control. For example, a skilled typist may intentionally perform a complex series of finger movements without having to separately intend each step in the sequence (i.e., each individual key press). That the individual key strokes aren't separately intended—i.e., that they happen automatically—doesn't disqualify them from being active. According to the account that I gave in Chapter 1, the individual key strokes are active despite being automatic because each key-stroke occurs under the control of the agent's intention—e.g., the intention to type a certain word or sentence. In just the same way, it is possible that attention unfolds in a simultaneously automatic, yet endogenously goal-directed manner.

The challenge the environmental salience view poses to actionism is therefore not that attention is sometimes automatic. Rather it is that attention is sometimes automatic in the way a knee-jerk reflex is—i.e., exogenous and stimulus-driven. That is why the first constraint on an adequate actionist response to the objection from attention capture—namely, that attention capture be shown to manifest an agential capacity—is not satisfied if the environmental salience view is true. At the same time, on the environmental salience view,

there is nothing inherently defective or unsuccessful about attention shifts that are purely salience-driven or exogenous. In being stimulus driven, automatic attention behaves exactly as it has the function of behaving. Consequently, if the environmental salience view is true, then the actionist's second constraint—namely, that cases of distracted capture be an unsuccessful or defective operation of attention—is also arguably not satisfied.

Having outlined the environmental salience view and the challenge it presents to actionist accounts, I'll now explain why I think we should reject the view. Grasping the reasons for rejecting the environmental salience view will, in turn, pave the way for my positive alternative.

2.4 Against the environmental salience view

Though the environmental salience view continues to exert an implicit influence on theorizing about attention, there are both empirical and conceptual grounds to reject it. The main empirical challenge concerns the contribution of factors other than objective salience that contribute to determining how attention automatically changes across time, including endogenous factors like task-relevance, affective significance, and reward history. As we'll see, appreciation of these other determinants of automatic attention ultimately undermines the explanatory informativeness of environmental salience in a theory of attention and points toward an alternative picture of automatic attentional control as anchored in the agent's motivational priorities.

The first line of evidence against the environmental salience view comes from research on gaze allocation in situations of natural vision—i.e., everyday tasks like making a sandwich or crossing a street. While it is important not to confuse visual attention with visual foveation (Posner 1980), in real-world scenarios visual attention and foveation are typically closely correlated. This is partly because attention to a location is probably necessary to program an eye movement to that location (e.g. Deubel 2008; 2014). And since (outside special laboratory contexts) eye movements tend not to be explicitly intended by the subject, eye-tracking in the context of natural vision tasks provides a rough window into the dynamics of automatic attention.¹

However, eye-tracking data in natural vision tasks (as opposed to in laboratory settings that present subjects with 2-D images on a computer monitor) offers little evidence

¹It might also be noted that the temporal profile of goal-directed eye movements corresponds more closely to that of automatic than to voluntary attention. Recall that it takes about 100-120 ms for automatic attention to reach a cued location and 300 ms for volitional attention to do so. The average latency of a goal-directed saccade—the time between the appearance of a cue and saccade initiation—is around 200-250 ms (see Carrasco 2011, 1490). Supposing that attention must already be at the target location in order to program the saccade, this suggests that these eye movements are programmed on the basis of automatic attention.

for the environmental salience view. For example, when reporting an experiment by Rothkopf et al. (2007) in which subjects had to navigate obstacles in an ordinary 3-D environment, Tatler et al. (2011, 4) note that subjects tended to look primarily at objects and to make only 15% of visual fixations to the visual background, whereas salience-based models predicted that humans should have made more than 70% of fixations to the background. Further, salience-based models have trouble explaining a common feature of human eye movements observed during performance of well-practiced tasks: the tendency to make anticipatory eye movements (saccades) to where a target object will be in the very near future. While playing a ball game, for example:

Saccades are launched to regions where the ball will arrive in the near future (Ballard & Hayhoe 2009; Land & McLeod, 2000). Crucially, at the time that the target is fixated, *there is nothing that visually distinguishes this location from the surrounding background of the scene*. Even without quantitative evaluation, it is clear that no image-based model could predict this behavior. Similar targeting of currently empty locations is seen in every-day tasks such as tea making (Land, Mennie, & Rusted, 1999) and sandwich making (Hayhoe, Shrivastava, Mruczek, & Pelz, 2003). When placing an object on the counter, people will look to the empty space where the object will be placed. (Tatler et al. 2011, 4, italics added)

In other words, when subjects automatically shift their gaze in anticipation of a fly ball or of a change that the subject will bring about by acting, there need be (and often will be) nothing physically salient about the location to which the eye moves. Yet, the attention deployed here is automatic. Tatler et al. don't deny salience plays a role in gaze allocation but do deny that it is anywhere as central to automatic attentional control as the environmental salience view suggests (Tatler et al. 2011, 10-11). In a similar vein, Henderson et al. (2007) and Henderson et al. (2009) offer evidence for what they call the "cognitive relevance framework" of overt visual attention (a.k.a. the "cognitive control hypothesis") against the "visual saliency hypothesis". They offer evidence that visual searches in the context of real-world scenes rather than 2-D computer displays are "fast and efficient" despite being very poorly predicted by environmental salience: "In the majority of trials, salient regions were not fixated" (850).

One might reply that while evidence from natural vision experiments may undermine the environmental salience view as a completely general account of automatic attention, the environmental salience view might still provide an explanation of *attention capture*. In particular, one might point out that it seems wrong to describe an expert's trained attention as being "captured" when she skillfully anticipates the upcoming location of the ball and moves her attention and eye to that location. Whatever troubles it might encounter when trying to explain the skilled attention experts manifest in natural vision

tasks, the environmental salience view could still provide a correct explanation of what happens when our attention operates incongruently with our goals.

In response, when we examine the empirical work that speaks most directly to involuntary attention capture (rather than automatic attention in general), the environmental salience view fares little better. For example, consider the orthodox suggestion that stimuli with abrupt visual onsets are especially potent at driving attention in a purely stimulus-driven, bottom-up manner. According to the environmental salience view, it is the intrinsic salience of abrupt onsets that explains why a direct/peripheral cue in a spatial cuing paradigm consistently captures subjects' attention (even when they intend to ignore it).

Against this, Folk et al. (1992) demonstrated that the apparently stimulus-driven character of subjects' attention to abrupt onset cues is probably an artifact of the experimental paradigm that researchers had been using to study attention. Using a variant of Posner's spatial cuing paradigm, Folk et al. sought to show that whether or not an abrupt onset cue succeeds at capturing a subject's attention critically depends on whether the target of the subject's current task also has an abrupt onset. For example, when subjects have the task of responding to an abrupt onset singleton, Folk et al. found that abrupt onset colourless cues captured attention, whereas non-abrupt onset coloured cues failed to do so. But when subjects were told instead to detect a red, non-abrupt onset singleton among a set of uncoloured distractors, they found the opposite pattern of results. In that case, an abrupt onset colourless cue *failed* to capture attention, whereas a non-abrupt onset red cue *succeeded* in doing so.

Folk et al. concluded that a stimulus captures attention only when it instantiates "a feature property that is critical to the performance of the task at hand" (Folk et al. 1992, 1032). In particular, stimuli capture attention only if they are represented in the agent's current task set or "control settings". To be clear, Folk et al. don't claim that such automatic shifts of attention are volitional or explicitly intended. Subjects may be fully aware that the cue that grabs their attention is a distraction and consequently they may intend to ignore it (e.g., in conditions where the cue is 100% invalid with respect to upcoming target location). Their hypothesis is, rather, that when a distractor *captures* attention, it does so only because it shares a perceptible feature with one's target. In a task like the spatial cuing paradigm, a peripheral cue captures attention only because it instantiates the visual feature that one is using as one's basis for target selection. If so, then attentional capture is "contingent" on task set:

The range of external events that can cause interrupts is determined by the internal settings of the software, which, unlike "hard-wired" interrupt systems, can be varied by changing the code. Once the code is compiled and the pro-

gram initiated, however, events that have these properties will be processed independently of the central processor (i.e., involuntarily). Similarly, in exogenous attention allocation, the control settings determining what properties will elicit shifts of attention can vary in accordance with task demands. Once the system is set and the processing sequence begun, however, the attentional response to properties compatible with current control settings is strongly involuntary. (Folk et al. 1992, 1041)

This is called the “Contingent Capture” or “Contingent Orienting Hypothesis”.² If it is right, then in adopting a certain task set, we bind ourselves in advance, like Odysseus to the mast of his ship, to having our attention automatically captured by a certain range of stimuli (cf. Elster 2000).

The Contingent Capture Hypothesis is controversial. However, even its critics (e.g., Belopolsky et al. 2010) do not attempt to explain Folk et al.’s key findings in terms of physical salience. Instead, what one finds is the notion of “salience” broadened to encompass other sources of attentional bias than that emphasized by, for example, Ulman and Koch (1985). One alternative to Folk et al.’s interpretation of the contingent capture effect highlights a confound with the subject’s recent history of attending to the target-congruent feature (through an entire block of trials). Given this recent history of attention to that feature, subjects might be primed to attend to the feature when it next appears (Belopolsky et al. 2010). If this is right, then the contingent capture effect might not directly depend on task set, but selection history. Does the latter reflect a purely “bottom-up” form of attention? One ground against an affirmative answer concerns the close relationship that may exist between selection and reward history. Awh et al. (2012)—two of whose authors overlap with Belopolsky et al. (2010)—propose:

The selection bias towards previously attended features or positions [as, e.g., defended by Belopolsky et al. (2010)] may be grounded in the reward that observers experience when they achieve their task goals (4).

Similarly, in their review of the effects of reward on inter-trial priming, Chelazzi et al. (2013) have suggested that:

[I]n standard experiments of this kind, where reward is not typically involved [i.e. explicitly administered by experimenters], the occurrence of negative priming may be dependent on some form of internal reward, acting in a similar way to external reward feedback. (63)

The suggestion Chelazzi et al. and Awh et al. each make is that the effect of recent (and not so recent) selection history on automatic capture may be the product of value

²A related view is the “Displaywide Contingent Orienting Hypothesis” (Gibson and Kelsey 1998). Burnham (2007) reviews the recent attention capture literature (including paradigms I have not discussed here) and concludes that, to date, there is no evidence for purely stimulus-driven attention in this literature.

learning mechanisms. The suggestion here is that subjects experience an implicit reward upon completing each trial, disposing attention to return to reward-associated stimuli when they next appear. This would make selection-history-based capture an instance of “value-driven attentional capture” (Anderson et al. 2011; Peck et al. 2009) likely rooted in “incentive salience” (Berridge 2004; Anderson and Yantis 2013). Concurrent to the discovery of value-driven automatic attention rooted in incentive salience has been the discovery of “affective” or “emotional” salience resulting in “emotional capture” (e.g., Fox et al. 2001; Vuilleumier 2005; Todd et al. 2012; Pessoa 2013), which likely interacts with reward history (Todd and Manigold 2018). Together with the possibility of task-contingent capture, this has led prominent figures in the field of attentional control to renounce the “top-down/bottom-up distinction of attention control” as a “failed theoretical dichotomy” (Awh et al. 2012). Given their dependence on the subject’s internal goal state and history, these non-intrinsic varieties of salience—not only task-contingent, but also emotional and incentive-related—are not well-understood as “bottom-up” or “stimulus-driven”.

Confronted with these difficulties for the original environmental salience view, one might propose amending the environmental salience view to accommodate other varieties of salience. One might suggest that all cases of attention capture are driven by the perception of environmental salience, but “environmental salience” should be conceived broadly to encompass not only physical or intrinsic salience, but salience in a sense that encompasses emotional, incentive, and task-contingent sources of attentional bias. Let’s call this the “liberal” environmental salience view of attention capture, as distinct from “conservative” environmental salience view that theoretically privileges physical salience.

However, it is hard to see how environmental salience could continue to play an informative role in a causal explanation of attention capture if understood liberally. On the liberal interpretation, stimulus salience will be extremely context- and response-dependent. It will depend on, among other things, the perceiver’s reward history with the stimulus-type, current level of arousal and motivational drive (e.g., hunger, thirst, etc.), task set, and possibly other factors too. There need be nothing in common among such stimuli other than being such that exposure to them will tend to trigger an automatic shift of spatial attention in some perceiver under some circumstances. Consequently, an explanation of attention capture framed in terms of its perceived salience will return us to the difficulties that the environmental salience view was originally introduced in order to overcome. In particular, to claim that stimuli are salient on the liberal construal amounts to a mere re-description that it is attention capturing (for a specific subject at a specific time), which is what we invoked salience in the first place to explain. So, whereas the original envi-

ronmental salience view was substantive but empirically false, the liberal salience view is potentially true, but uninformative.

The considerations discussed in this section suggest that we should reject the environmental salience view. We need a different account of the mental state underlying automatic shifts of attention. We could continue to call this mental state “salience”, and it may be that some empirical researchers intend to use the term in this broad manner. However, given the close associations of the term “salience” with the specific commitments of the environmental salience view, I will mostly avoid using the term henceforth. I will speak instead of the mental state responsible for the control of automatic attention. In the next section, I offer an account of this mental state that grounds it in the agent’s priorities, and I show how this helps answer the objection from attention capture against actionism.

2.5 The priority view of automatic attention

This section develops the positive view of automatic attention that I want to propose. I call this view “the priority view”. The priority view has the following main constituents. It understands automatic attention, including various cases of captured attention, in relation to attention’s aim of preventing distraction (Chapter 1 see §1.5). The priority view retains this claim but it casts it in a more positive light. It claims that the aim of automatic attention is to prioritize. “Prioritization” is the activity of aligning information-processing with the agent’s goals (or “agential priorities”). When all goes well, this process results in states of processing in which information that would otherwise interfere with successful action is suppressed and goal-relevant information is selectively enhanced. Distraction is averted and relevant information is prioritized. Importantly for answering the objection from attention capture, it is possible, on the priority view, for an agent to attend but to fail at prioritizing. The priority view therefore does not claim that attention *just is* prioritization, but rather that prioritization is the proper functional role or aim of attention. When we are attending well, we are prioritizing.

Before proceeding, I want to clarify how the priority view relates to relevant empirical work. Recall the empirical evidence raised against the environmental salience view in the previous section. That evidence suggests that environmental salience is not the only, or even a primary, causal factor in determining how one automatically attends. In addition to environmental salience, multiple endogenous factors appear to contribute as well, including task set, reward history, and affective significance. In recognition of these findings, psychologists of attention have increasingly begun to adopt the expression “attentional” or “integrated priority map” (or simply “priority map”) in place of the more traditional “salience map” when referring to the representational structure thought to be

responsible for directing automatic visuospatial attention (e.g., Fecteau & Munoz 2006; Serences & Yantis 2006; Bisley & Goldberg 2010; Baluch & Itti 2011; Awh et al. 2012; Todd and Manaligod 2018). We are told that the priority map integrates both bottom-up information arising from the sensory receptors (including about environmental salience) with top-down information in order to construct a topographical ranking of each location in the scene according to the overall “priority” (Fecteau and Munoz 2006; Baluch and Itti 2011) or “value” (Gottlieb 2012) of the information at that location given current behavioural context. Beyond this minimal characterization, researchers rarely explain what theoretical significance we should attach to this terminological shift, except to remark that it reflects an increased emphasis on endogenous factors like “relevance” in determining the content of this map. The priority view is my own attempt to identify the broader significance of this shift for understanding the nature and agential status of automatic attention. So, while the priority view is anchored in recent empirical work on attention, I am going beyond what psychologists themselves say about the matter.

This is a long section, so here’s my plan for it. I begin in §2.5.1 by elaborating on the priority view and its main differences from the environmental salience view with respect to the agential status of automatic attention. In §2.5.2, I explain how the priority view claims to accommodate cases in which the aim of prioritization is intuitively unsatisfied, including cases of distracted attention capture. I close in §2.5.3 by drawing a comparison between my account and two other extant proposals in the philosophical literature.

2.5.1 Agential priorities, acts of prioritization, and the integrated priority map

Let’s begin with an example to illustrate the main idea behind the priority view. Consider the visual attention of a well-motivated, expert squash player who is visually tracking the squash ball as it leaves her opponent’s racquet. The squash player is motivated by the goal of successfully returning the shot and ultimately of winning the match. Because she is thus motivated, she ignores what would otherwise distract her—e.g., the bright sheen of her opponent’s watch or the loud conversation occurring outside the court—and instead focuses on the parts of the environment that are relevant to attaining her goals—e.g., she might predictively saccade to the upcoming location of the ball (see §2.4). Suppose that she succeeds: her visual processing of the ball’s upcoming location is selectively enhanced and visual processing of the location of various distractors is suppressed. By reorienting her attention automatically to where the ball is soon to be (and away from, for example, the watch or the loud events beyond the court), she attends to what’s currently most relevant to her. By aligning her perceptuomotor processing with her goals in this way, the

squash player qualifies as engaged in a kind of prioritization. According to the priority view, such prioritization is the constitutive aim of automatic attention.

Importantly, the notion of “priority” enters the priority view at multiple junctures. Most fundamentally, we must distinguish the agent’s *priorities* (or “priority system”) from the agent’s acts of *prioritization*. I’ll begin by explaining the former and how it relates to the psychological structure of attention, as this was explained in Chapter 1 (§1.3).

By “agential priorities”, I mean the agent’s total or overall motivational state at a time. This includes the integrated goal hierarchy provided by an agent’s intentions and plans. In the above example, these consist in the player’s intention to return the ball and ultimately to win the match. However, as I will understand the notion of an “agential priority”, an agent’s intentions are only one component of her priorities. This is because intentions do not exhaust the agent’s total motivation. To illustrate the importance of total motivation, consider the difference between someone who is strongly motivated to win a squash match and someone who is less motivated to win. For the first person, winning is clearly a priority. About the second individual, though, we might say that while she also had goal of winning, her “heart wasn’t in it”. For the latter individual, winning isn’t, at least in the moment, as much of a priority. There is a difference in the overall motivational strength with which the agents hold the goal of winning—a difference that seems to directly impact this goal’s status as a priority.

What exactly this difference in overall motivational strength in the above example amounts to isn’t totally clear. However, it likely goes beyond the contribution of any single attitude. As Shepherd (2017b) observes in a recent discussion of motivational strength,

The dispositional basis of one’s overall motivation to A will consist of the dispositional properties of a range of accessible motivational mental states and processes relevant to one’s A-ing. These may include emotions (Scarantino & Neilsen 2015), desires (Mele 2003), intentions (Mele 1992), as well as modificational work done by relevant beliefs, perceptual states, states of imagination, and more. (Shepherd 2017b, 262; cf. Mele 2003, 173).

Building on this point, I propose to understand the agent’s current priorities as determined by the integration of a range of motivational systems. This includes not only conscious intentions but also implicit and automatic affective and reward-related biases, levels of arousal, and potentially more. To emphasize the wide range of contributing factors, I’ll also speak in this context of the agent’s “priority system”.

With the notion of the agent’s priority system in hand, we must revisit the account of attention that I sketched in Chapter 1 (§1.3). There, I explained that attending is a biasing process composed of two mental states: an attentional source state and an attentional resultant state. The former is the source of attentional bias, and the latter is the result

of the process on lower level states of sensory and sensorimotor processing. In Chapter 1, I operated with the simplifying assumption that the attentional source state during action is the agent's intention in acting, which I related to a task set or intentional set. However, we must amend this to accommodate the broader notion of agential priorities. In particular, I want to suggest that where an agent's intention contributes a source of automatic attention, it does so in virtue of the place that it occupies within the agent's larger priority system. At least in the case of automatic attention, it is the agent's overall priority system or total motivational state that serves as the source of attentional bias. I will continue to refer in the singular to the attentional source state for an episode of automatic attention. But we should bear in mind that this source state is an aggregate of motivational factors.³

According to the priority view, when agents automatically guide their conduct, they do so according to the standards of success set by their agential priorities. Attentional prioritization provides the means whereby agents exercise such automatic guidance. It is critical to grasp that, on the priority view, what makes information-processing an instance of attentional prioritization is the role of the agent's priorities in directing it. For example, when our well-motivated squash player reorients her attention to the ball's upcoming location, she counts as having prioritized that location only if it was her priority system that guided the reorienting. Otherwise, as I'll discuss more fully in §2.5.2, she would not qualify as having prioritized at all.

To consolidate the view, it is instructive to contrast the priority view's conception of "attentional prioritization" with how the same expression is apt to be interpreted on the environmental salience view. If by "attentional prioritization", one merely meant the selectively enhanced processing of some information to the exclusion of other information, then the claim that the proper functional role of attention is to psychologically prioritize information is fully consistent with the environmental salience view. After all, the environmental salience view fully accepts that perception of environmental salience triggers selective processing of the salient location. However, by understanding "prioritization" in this way, one would be severing any essential connection between prioritization and the agent's motivational priorities. Consequently, the priority view does not understand the function of attention in that way. In formulating the role of attention, the priority view instead draws on an understanding of "prioritization" according to which prioritizing is a process whose standards of success are constitutively determined by the agent's prior-

³Further, the neural correlates of these diverse states may be more physically distributed than sometimes suggested. Attentional source regions include not only the frontal "dorsal" and "ventral attentional networks" Corbetta and Shulman 2002; Shulman and Corbetta 2012, but also portions of the amygdala for affective biases (Allport 2011, 40; Pessoa 2013, Ch. 2) and possibly also sources of motivational bias in the basal ganglia (Cisek 2007).

ity system. Specifically, the priority view means by “prioritization” the act of bringing one’s states of information-processing into alignment with one’s agential priorities. Understood in this way, while it may be possible for selectively enhanced processing to be triggered purely “bottom-up” on the basis of perceived salience, it is not possible for an act of attentional prioritization to occur that way. The latter act must, instead, issue from the agent’s priorities.

With this, we are in a position to grasp why the priority view qualifies as form of actionism about attention. The priority view regards *all* episodes of automatic attention as motivated occurrences and, in this sense, as emanating from the agent. Whereas the environmental salience view depicted an automatic attentional shift as a brute, reflexive response triggered by the perception of salience, the priority view understands an automatic shift of attention as a motivated response to the detection of goal-relevant or “high priority” information—i.e., information that *merits* prioritization. This includes instances of what we would intuitively characterize as “captured” attention. After all, sometimes attending to what is of highest priority to you requires automatically interrupting what you’re doing in order to attend to a more urgent concern. For example, while our squash player is preparing, preemptively, to move to the front of the court in anticipation of a boast from her opponent, suddenly her attention becomes captured by an unexpected feature of her opponent’s follow-through which has the potential to drive the ball in a direction opposite to the anticipated one. The attention that she was paying in expectation of a boast has suddenly been interrupted by an unexpected, but highly relevant visual cue. In these cases, an agent’s attention and behaviour may be fully coherent and may even manifest significant skill. Thus, even when captured, on the priority view, attention is directed from *within* by the agent’s priorities rather than from *without* by environmental salience. Because it regards *all* episodes of automatic attention as possessing a motivational source in the agent’s priorities, the priority view meets the first of our constraints on an adequate response to the objection from attention capture. I will examine more complex cases of capture in §2.5.2.

I’ve been relating the priority view of automatic attention to my thesis in Chapter 1 that attending is a process anchored to the agent’s goals and aimed at coping with distractions. According to the priority view, this is the theoretical context within which we must understand the “integrated priority map” that psychologists working on the control of visual attention posit. For the remainder of this subsection, I comment on what significance we should attribute to this map on the priority view.

Among psychologists, the assumption that there is a literal topographical map in the primate brain that drives subsequent states of visuospatial attention in virtue of its content remains fairly entrenched. Much research is devoted to investigating the proper-

ties of this map—e.g., determining which factors (both exogenous and endogenous) contribute at which time scales to its rankings. However, while I do not wish to downplay the importance of the priority map, it is possible that these assumptions are partly a holdover from the environmental salience view. Once we have rejected the assumption that automatic attention is driven exogenously by the *perception* of salience, we may wish to reconsider the assumption that automatic attention is controlled by a *map*. Rejecting this assumption is consistent with retaining the empirically well-established construct of the priority map, though we may wish to reinterpret its exact place in attention.

In particular, once we reject the environmental salience view and embrace the priority view, the natural conclusion to draw is that the actual source state for attention is not strictly the priority map itself. Rather, the source state is the agent's own priorities: the set of systems that are contributing *inputs* to the priority map. Once these disparate inputs have been integrated into a unified visuospatial representation of the scene in the form of a priority ranking, one might plausibly suggest that a form of attention is thereby instantiated in the state of the map. The agent's motivational priorities have been brought to bear on available visual information and resulted in the ranking of visual information according to its relevance to priority attainment. And that *just is*, on the priority view, a form of attention. Importantly, the agent's act of attention may be *incomplete* or *in progress*, and the priority map certainly has an important coordinative role to play in how that attentional activity develops. This makes the priority map an early, though functionally critical stage in a temporally extended process of attention. But this is consistent with maintaining that the state of the priority map already instantiates a form of visual selective attention.

Moreover, although it has been the focus of a large amount of work on the control of visual attention, the contribution of the priority map might turn out to be relatively narrow from the perspective of guiding bodily action as a whole. In particular, the brain areas researchers have identified as the basis of the priority map in primates—the well-studied regions of posterior parietal cortex, including the monkey homologue lateral intraparietal cortex (LIP)—are significantly overlapping with (if not just the same as) areas that program rapid eye movements. Furthermore, it is known that this map has a spatial frame of reference and functional interconnections (e.g., to the frontal eye fields) tailored to this specific role. There are many ongoing controversies about the different functions carried out in LIP and functionally connected areas like the frontal eye fields (Armstrong 2011). But if one is interested in how visual attention is used to guide other aspects of bodily action, like reaching out one's arm, then other areas with very different spatial format and functional connections could turn out to be more relevant (e.g., medial intraparietal cortex or MIP). While this does not preclude the priority map in LIP and connected areas from playing a privileged role in orchestrating these other resources, it raises the possibility

that the priority map is one specialized resource among many others that is motivationally harnessed in attending. To this, we can add that there are forms of attention other than spatial attention, including object-based and feature-based, which the priority map won't be able to account for given that it is location-based. This is further reason to doubt that it will play an essential role to all cases of visual attention.

None of this is meant to deny or downplay the importance of the priority map in mediating at least one very central form of visual attention: attention of a kind that guides overt eye movement. And although I've proposed that the state of the map already manifests a form of selective visual attention, I do not deny that the priority map, in turn, plays an important role in biasing processing within, e.g., the ventral visual stream (e.g., feature maps in V1 and V4) and in motor areas like FEF toward the highest ranked locations in the map. In doing so, it serves to direct selectively enhanced perceptual and motor processing in order to bring that processing into alignment with its own priority rankings. Consequently, researchers are right to claim that the priority map genuinely *directs* visual attention. The points I've been emphasizing are simply that: (i) the priority map's role in directing visual attention is plausibly an intermediate stage within a larger attentional process that has its ultimate source in the agent's priority system; (ii) that there could be other structures in the brain that perform a similar mediating role in the flow of visual attention. The priority view has additional consequences for how we should conceive the role and content of the priority map. As these are not immediately relevant to the points I wish to make for the remainder of the chapter, however, I reserve discussion of these further implications for an appendix to this chapter. I turn now to explain how the priority view hopes to accommodate cases in which the aim of prioritization is intuitively unsatisfied.

2.5.2 Failures of prioritization

According to the priority view, all cases of automatic attention issue from the agent's priority system. By casting even cases of captured attention in these terms, the priority view claims to satisfy the first constraint on an actionist reply to the objection from attention capture. In this section, I show how the priority view meets the second constraint. The second constraint is to show why, in some cases of capture, attention operates poorly or imperfectly. Specifically, in some cases of attention capture, the agent's attention is distracted and, as a result, behavioural coherence decreases. Since the function of attention is to compensate for distraction and, in so doing, to promote behavioural coherence (see §1.5), it follows that attention fails to fulfill its role in these cases. The priority view recognizes multiple paths to prioritization failure. Charting these paths will allow us to grasp

the limitations of automatic attention as a form of control and the different sources of distraction that skilled agents must face in acting, ranging from the purely internal to the purely external.

The first way prioritization failure can occur is if the agent lacks a fully coherent system of priorities. Intuitively, this is a form of distraction whose source is within the agent. To develop this point, I need to introduce the concept of “motivational interference”. With this, I refer to the processing interference that arises when distinct motivational aspects of the agent’s priority system compete for control of attention. Conflicting priorities present an obvious barrier to successful prioritizing. It is important to observe, however, that motivational conflicts do not undermine attention. Automatic attention is simply the top-down biasing of perceptual and perceptuomotor processing by the agent’s priority system. And a biasing process can occur whether or not the source of bias is internally coherent. For example, if two individuals decide to direct traffic at an intersection but do little to coordinate their activities, they will be biasing traffic flow even if they are not doing so coherently (as measured by the likelihood of collision). Similarly, it is possible for the motivational elements comprising the agent’s priority system to generate states of processing that, while motivationally modulated and therefore attentional, are nevertheless not fully coherent. Here, the agent’s priority system continues to attentionally bias perceptual and perceptuomotor processing, but the coherence of attention (and ultimately of behaviour) is impaired because of motivational interference. It is distracted and thus defective attention.

On the priority view, many central examples of captured attention turn out to be reflect motivational interference. To illustrate, return to the example of experts “choking under pressure”—i.e., performing significantly below their skill level because of high anxiety about failure. As discussed in Chapter 1 (§1.4.2.3), according to one influential account (“attentional control theory”), experts control their skilled performances, in part, by deploying “goal-directed” attention when guiding a task’s execution. Further, according to this account, choking occurs when the expert’s attention automatically shifts away from task-relevant information toward task-irrelevant information—i.e., when they become distracted. (For this reason, the theory is also called a “distraction” theory of choking). Thus, when the expert golfer chokes, her attention is diverted *from* the ball and her swing, say, *to* distractors in her environment or worries about failure. It is this automatic attention shift that is claimed to underlie the expert’s sudden control loss and failed putt (Vine et al. 2013).

Unsurprisingly, the attentional control theory links this automatic attention shift specifically to anxiety. In making this suggestion, proponents of the account draw on the well-established finding that anxiety is a partial determinant of automatic attention.

For example, with respect to tasks that require detecting and responding to threat-related stimuli, anxious individuals are at a clear advantage over non-anxious subjects. In visual search tasks that require one to detect an angry face in a crowded display of emotionally neutral faces, anxious individuals are significantly quicker than non-anxious individuals to notice and respond to the angry face (Byrne and Eysenck 1995). When the search target is, instead, a happy face in a crowd of angry faces, the opposite pattern is found: search time for the anxious subjects is impaired relative to that of non-anxious subjects. On the basis of these and related findings, proponents of attentional control theory therefore propose that choking consists in an automatic anxiety-induced attention shift away from task-relevant information and toward threat-related stimuli or thoughts.

The priority view regards choking as an example of attention capture based in motivational interference. On the one hand, the golfer's intention to swing motivates devoting processing resources to the ball and to key aspects of her swing. On the other hand, her anxiety motives devoting processing resources to threat-related information (e.g., her anxious thoughts or the threatening-looking audience members) to the *exclusion* of the ball and swing execution. Here, processing that is relevant relative to one motivational factor is irrelevant relative to the other, and so interference arises. While this is how the priority view interprets these claims, it should be noted that this is not how proponents of attentional control theory *typically express* their position. More commonly, proponents continue to frame their account, as described in the previous paragraph, in terms of a traditional distinction between a "bottom-up" and "salience-driven" form of attention and a "top-down", "goal-directed" form of attention (e.g., Eysenck et al. 2007, 343). On this formulation, choking consists in a shift from goal-directed, top-down attention to salience-driven, bottom-up attention. The priority view rejects this aspect of the account. After all, choking is *not* salience-driven if by "salience" one means environmental salience. As Eysenck et al. themselves emphasize, attention during choking is directed by *anxiety*, which they aptly explain as "an aversive emotional and motivational state" concerned with the presence of threat (Eysenck et al. 2007, 336). What proponents of the theory *should* say is that choking results from motivational interference between one aspect of the priority system and another. Anxiety-induced capture has less in common with a brute reflex than it does with, for example, unskillfully attempting to do two things at once.⁴

Motivational interference tends to undermine an agent's intentional control. Although

⁴DeCaro et al. (2011) consequently come closer to adequately formulating the attentional control theory when they summarize its core claim as that "pressure essentially creates a dual-task environment in which situation-related worries compete with the attention needed to execute the task at hand" (391). However, this formulation isn't quite right either. Specifically, it is wrong to suggest that the expert's anxiety provides her with a *task* in addition to that of executing the skilled performance (as DeCaro et al. implicitly acknowledge by ending their sentence by speaking of "the" task at hand). What DeCaro et al. mean is that anxiety *motivationally interferes* with the subject's intention to execute a task (for control of attention).

it is possible that the anxious golfer will execute a flawless performance, her attention is more likely to move to task-irrelevant stimuli when focus on the ball and her swing is critical. If she nevertheless executes a perfect shot despite being anxious, luck will have been more of a factor in her success than if her attention had been more fully engaged on the task at hand. Similar points will likely apply to acting with weak overall motivation: when you are tired or bored, attention will tend to be riveted by more personally interesting topics, resulting in task-neglect (cf. Shepherd 2017b). And when trying to break a bad habit, “value-driven” attentional biases based in reward history can interfere with successfully acting on one’s intention (e.g., Berridge 2004; Robinson and Berridge 2008; Anderson et al. 2011; Anderson and Yantis 2013). More generally, if attention exists to achieve behavioural coherence through the avoidance of distraction, this role will tend to be sabotaged by motivational interference. It is difficult to attend well—i.e., to prioritize—if one’s priorities are not coherent. Motivational conflicts tend to undo what attention exists, in the first place, to achieve.

Given the many ways motivational interference can undermine attentional prioritization, one might wonder how agents cope with distractions that have a motivational source. To answer that question, I think we would need to go beyond automatic attention. Automatic attention is at the whim of total motivation or the agent’s overall priority system, including arousal levels (e.g., anxiety vs. boredom) and affective tendencies linked with habit and personality. One might conjecture that a more volitional and effortful form of attention developed, in part, to aid with the motivational interference that can plague automatic attention. In particular, a more volitional and effortful variety of attention might be used to bias the outcome of motivational competitions for attention in a certain direction. Though there will likely be limits to the efficacy of effort, there is evidence that the likelihood of choking under pressure decreases with increased effort (Eysenck et al. 2007). This might consist in volitionally directing effort to how one attends in light of one’s awareness of one’s attentional biases. It seems to imply making a decision about how one will attend, whereas automatic attention does not depend on deciding how one will attend (see Chapter 1 §1.3). Automatic attention may be sufficient to act successfully when one’s priorities are coherent; otherwise, effortful attention might be required.

I’ve been focusing till now on examples of distracted attention that have their basis in the attentional source state—i.e., the priority system. One might ask whether all cases of distracted attention can be explained in terms of motivational competitions within the agent. In response, I think that would be too simplistic. It is true that, on the priority view, many *more* cases of attention capture are apt to be explainable in this manner than has traditionally been assumed. However, we should allow that attention can fail in its aim

without there being any defect or incoherence in the agent's priorities. For the remainder of this section, I want to highlight two other ways prioritization failures could arise on the priority view. The first concerns task-contingent capture (Folk et al. 1992). The second concerns the influence of physical salience.

To begin, recall the putative phenomenon of task-contingent capture (Folk et al. 1992). These are cases in which a peripheral cue captures attention because it possesses a feature in common with one's target. For example, something red might automatically summon attention if one's task requires finding something red. This is found to occur even if the subject knows that a cue will appear before the target and is confident that the cue is merely a distraction that should be ignored. Nonetheless, Folk et al. argue that target-congruent cues reliably impair target detection by capturing attention. Let's suppose such cases occur and that Folk et al.'s explanation of them is the correct one.

What makes this case seem to stand apart from those considered earlier (e.g., being drawn to attend to stimuli associated with threat or with past reward) is that, in this case, there is a strong case for saying that the cue is wholly irrelevant to the agent. Since the subject does ignore the cue when it is task-incongruent, the only plausible basis of significance to the subject is its connection to her current task. And yet, the cue is task-irrelevant and understood as such by the subject. It seems that the agent's priorities could be fully coherent in a case like this one. The subject could be single-mindedly focused on the task at hand and motivationally well-balanced. And yet distraction arises in the transition from the agent's priority system to the state of the priority map. If we understand the priority map as a representation of the relevance of incoming visual information to the agent given her priorities, then the priority map misrepresents locations that are in fact of low priority (namely, the location of the cue) to the subject as being of high priority. Under the guidance of the nonveridical priority map, what is in fact a distractor receives selectively enhanced visual processing and performance suffers as a result. The agent's attention is captured by a distractor. Here, distraction arises not because of interference within the priority system, but because of a mismatch between the agent's priorities, on the one hand, and the state of her priority map (and, ultimately, her completed state of attention), on the other. In principle, nothing precludes other forms of misrepresentation from entering at this stage as well.

Whatever the empirical fact of the matter about task-contingent capture, the possibility of such cases should not surprise us. This is so especially when we recall that we are dealing with an automatic and pre-attentive process. The calculations that go into determining a new state of the priority map (i.e., a new priority ranking) are occurring outside the focus of attention (since they are responsible for determining what will be focused on next). Consequently, one might well expect these processes to be less flexibly

sensitive to the agent's explicit goals than processes that are possible with the benefit of focused attention. Further, we might expect the limits of the flexibility of those automatic processes to come to light under special laboratory conditions similar to those Folk et al. designed. The process whereby the agent's psychological processing becomes aligned with her motivational priorities is, unsurprisingly, fallible.

I've suggested that task-contingent capture may present a case in which discrepancy arises between the agent's priority system and the state of her integrated priority map (and ultimately of her completed state of attention). In such a case, the map represents a location as having high priority to the agent which, given the agent's actual priorities, is of very low priority. Whereas this type of failure arises in the transition from the agent's priorities to the state of her priority map, the last type of case I will consider points to a role for purely bottom-up influence in undermining attentional prioritization. These are failures of prioritization that arise after the priority map has been formed, but before the priority map has successfully biased other areas. If we understand the priority map as an initial stage in an attentional process, then this sort of failure would arise after attending is already underway. In this context, one possibility that deserves consideration is whether environmental salience could exogenously influence the course of attention at this point—i.e., after the priority map has generated a priority ranking. I've already allowed that environmental salience may bias a priority map's rankings. What I'm entertaining now is the possibility that environmental salience also exerts a biasing influence on the course of attention that is unmediated by the priority map. If perceiving a physically salient distractor could cause the unfolding process of attention to deviate from the priority map's ranking, this could, in turn, occur in either of two ways. First, it could be that salience merely *interferes* with attending so that it deviates from its priority ranking. We can think here of the buzzing fly tugging on one's attention as one tries to work. Second, it could be that the salient distractor utterly *overwhelms* the priority map's biasing and so fully "captures" perceptual processing (cf. Buehler 2014). In this second case, whatever bias the priority map had (or would have had) is nullified. We might think here of a loud crash or explosion.

The priority view implies a very different view of the two cases with respect to the presence of attention. On the plausible assumption that all actions are *influenced* or *biased* externally, the first scenario would not undermine the status of attending as an exercise of an agential capacity. The squash player's swing, for example, is influenced by the physical properties of the racquet itself (e.g., its weight), but she continues to actively guide its progress. That an action's course is *biased* exogenously doesn't entail that it is *determined* exogenously—i.e., stimulus-driven. The latter requires not only the presence of exogenous factors, but the absence of endogenous ones. In the case where environmen-

tal salience merely biases attention from below, the agent's priorities remain engaged in guiding attention. There is merely interference from a foreign element, which diminishes but need not eliminate agential control. In the second scenario, however, matters are different. If salience not only biased but *overwhelmed* the influence of the agent's priorities, then on the priority view this would not be attention at all, but pure distraction. A purely exogenously induced change in perceptual processing is not a change in *attention* (see also Chapter 4).

It should be noted that while nothing on the priority view precludes salience from overwhelming the influence of the priority map, scenarios like automatically reorienting to a loud crash involve a confound that is important in the present context. Physical salience, on its own, may have *some* automatic biasing effect on attention (though the matter is unclear, see Burnham 2007). But what exerts a *strong* automatic bias on our attention, and has the potential to completely interrupt our focus, is the emotional *valence* of a stimulus. Novel or infrequent information can be highly emotionally arousing. And it is novel or infrequent stimuli that appear to engage the brain areas associated with the "interrupt" of attentional focus rather than physical salience *per se* (see Shulman and Corbetta 2012, 117-118).⁵ Physical salience and emotional valence are hard to disentangle in examples like a loud, unexpected crash because these events are highly emotionally arousing. They startle us and cause fear. Since I am including emotional systems as a motivational constituent of agential priority, cases like the loud, unexpected crash are already accounted for on the priority view.

I've considered the main ways attention can potentially fail to achieve its proper functional role or prioritization. In each case, we find acts of attention that result in a state of at least partial distraction—i.e., in failures to prioritize what's relevant given one's agential priorities. This delivers the result I've suggested we should want from an actionist conception of attention: it depicts attention as essentially the exercise of an agential capacity. Yet it accommodates the observation that attention can be captured in a way which results in distraction or failure to prioritize. In such cases, attention continues to be governed by the aim of prioritization, but, for a variety of possible reasons fails to achieve that aim and so is flawed attention. Recognition of the category of defective attention is absent from

⁵In particular, Shulman and Corbetta discuss a study by Indovina and Macaluso (2007) showing that "irrelevant stimuli of high sensory salience that occur as a regular part of an ongoing task, in which other stimuli have been designated as behaviorally relevant, do not activate the ventral network" (where the ventral network is the one that, on Shulman and Corbetta's influential framework, underlies the capture (or "interrupt") of focus from a task to something more behaviourally significant) (Shulman and Corbetta 2012, 118). Shulman and Corbetta quite explicitly link this and related findings to the known anatomical and likely evolutionary ties between the right-lateralized attentional ventral system and the known seats of emotional arousal (which, they note, also tend to be strongly right-lateralized). The central thrust of Shulman and Corbetta's discussion is not that the ventral network is "stimulus-driven" (as they themselves tend to put the point), but rather than it is primarily *emotionally* driven.

extant actionist proposals, including proposals that are otherwise strikingly different, as I'll now explain.

2.5.3 An alternative to “rigorism” and “permissivism”

In this section, I contrast the priority view of automatic attention with two more extreme views of attention: what I'll call “rigorism” and “permissivism”. I'll argue that the priority view avoids shortcomings confronting each position.

According to rigorists, attention is *essentially* an absence of distraction. Like proponents of the priority view, they maintain that the nature of attention is to be understood in relation to the agent's priorities or goals. They further claim that wherever one is distracted by what one understands to be low-priority or goal-irrelevant information, one thereby fails to be in a state of attention. For example, if subjects understand that a direct cue in Posner's spatial cuing paradigm is a mere distraction, rigorists are likely to deny that the allocation of selectively enhanced visual processing resources toward the cued location realizes a form of attention at all. According to the rigorist, matters would be different if, once one had been distracted by the cue, one took up a new task whose performance was understood to involve selectively enhanced visual processing of the cued location. In that case, we could then, in rigorist terms, speak of the agent's *attention* as having been captured by the cue. But suppose the agent takes up no new task. In that case, rigorists must regard the subject's impaired performance in the invalid condition not as a consequence of the subject visually attending to the cued location, but of the agent *merely* being distracted by the cue. Like the priority view, rigorists emphasize that an episode of attention should be assessed relative to the standards of success established by the agent's goals or priorities. But unlike the priority view, rigorists are led to reject the possibility of “distracted attention” as an oxymoron. Mole's (2011) account can be read as suggesting this verdict about the effects of a direct cue in Posner's spatial cuing paradigm.

Permissivists, by contrast, happily accept the idea of distracted attention. They can therefore allow that the agent counts as visually attending to the directly cued location in Posner's spatial cuing paradigm, even though they understand the cue to be a mere distraction. The permissivist goes further, however, and claims that there is nothing sub-optimal or defective about attention in such cases. At least by the standards of success imposed by attention itself, one attends fully successfully in being distracted. According to Watzl (2017), for example, attending consists in the mental activity of structuring one's mental states into relations of relative priority: what Watzl calls “priority structures”. In cases of capture by a direct cue, Watzl claims that the agent's structuring activity is guided

by “psychological salience”, which he takes to be an imperatival aspect of perceptual representational content. For example, on Watzl’s view, the subject’s perception of a direct cue in a spatial cuing paradigm includes, as part of its intentional content, the imperative <prioritize this!> or <put this perceptual state at top of priority!>. That one’s perception of the cue is a distraction from one’s task is, on Watzl’s view, orthogonal to the question of whether one qualifies as having attentionally prioritized it. On Watzl’s view, one *prioritizes* the distractor in the fullest sense of that term.

Rigorism and permissivism are each open to criticism from the other side. To begin, permissivists will criticize rigorists for being out of step with the science of attention, as well as a central fragment of common sense thought about attention, when they suggest that things we understand to be distractions from our goals can never be objects of our attention. When a direct cue flashes, for example, we are not only distracted by it; our attention is captured. And this can be so, even if we have at no point abandoned our original task. To deny this is to be simply at odds with the science of visual attention and folk intuition.⁶

Rigorists, in turn, can criticize permissivists for draining attentional prioritization of its normative dimension by erasing any essential link between this activity and the agent’s actual priorities. Permissivists are thus in a position similar to that of the environmental salience view. Specifically, permissivists seem to inadequately accommodate the intuition that often it is our being distracted that *constitutes* our failing to prioritize. For example, suppose Jane has the priority in the morning of getting ready to leave her apartment on time. In that case, Jane’s prioritizing will consist in, among other things, avoiding getting absorbed in an engaging conversation with her roommate about current affairs or browsing the internet while sipping leisurely on her coffee. Likewise, she will have failed to prioritize in the relevant sense if she becomes involved in one of these activities without having taken the steps required to leave her apartment on time. The permissivist

⁶Wu (personal correspondence) has suggested an interesting revisionary interpretation of the facilitation effects of the direct cue in Posner’s spatial cuing paradigm to which a rigorist might appeal in order to deny that subjects shift attention to the cued location. He suggests that perhaps the direct cue merely *primes* one to attend to the cued location, such that were one’s target to appear there, one would be quick to attend to it (i.e., on Wu’s view, select it for action). However, I do not think this proposal sits well with how researchers working on attention capture conceive the relation between attention and priming (e.g. Chelazzi et al. 2013). Very roughly, a running theme of this work is that the subject’s history of attentional selection history with a certain target-type (in roughly Wu’s sense of having selected the target for task-response) in earlier experimental trials primes the subject to select stimuli of that type again when they later appear as *cues* in later trials. That is, if one has selected a target for action in the past (likely with some accompanying reward), and if the cue in a subsequent trial is of the same stimulus-type as the target that was earlier selected for action, then the subject is already primed to select the *cue* when it appears in the later trial. Wu’s proposal requires that the cue prime one to attend to the target, but it seems that one is at least sometimes already primed to attend to the cue at the start of the trial precisely because of one’s history of selection for action in Wu’s sense. So, this alternative interpretation of the effect of the direct cue seems not to provide a successful reply to the objection I have attributed to the permissivist above.

might reply that, on an imperativist account of “salience”, Jane possesses a reason to prioritize such things as reading the news in the morning, even if doing so takes her further from achieving her expressed goals. To this, the rigorist can plausibly respond that merely being *told* to prioritize something does not give one any reason to do so (even if, as on Watzl’s view, it is your perceptual state issuing the imperative) (cf. Aydede and Fulkerson 2019). To have a reason to prioritize something requires that it bear some connection to one’s own priorities.

The priority view of automatic attention I’ve sketched in this chapter presents us with an alternative to permissivism and rigorism that avoids the shortcomings of each. The alternative is, briefly, to recognize two standards as relevant to assessing an episode of attention and to maintain that cases of distracted attention capture satisfy only one of the two standards. The first standard establishes the presence or absence of attention. It concerns whether or not the agent’s motivational states are directing information-processing. If they are, then attention is instantiated. The second standard establishes whether or not attention fulfills its proper functional role. This concerns whether or not the motivationally modulated information-processing successfully removes states of distraction. If the latter condition is met, then not only does one *attend* in virtue of motivationally biased information-processing, but one *prioritizes* in the full-fledged sense of resisting distractions in service of one’s priorities. In this way, it is possible to attend distractedly (contra the rigorist), but distracted attention is deficient as attention (contra the permissivist). Rigorists mistake the second standard for an existence condition rather than a success condition, whereas permissivists (like the environmental salience view before it) simply fail to recognize the second standard at all. The permissivist is therefore right to complain that the rigorist makes attention too difficult. Attending is easy, but attending *properly* is difficult.

2.6 Conclusion

In this chapter, I have argued that every episode of attention consists in the exercise of an agential capacity. Specifically, it consists in the exercise of a capacity to prioritize what is important to the agent in light of her overall system of priorities. Combining this with the less controversial claim that non-automatic deployments of attention are also exercises of an agential capacity, we reach the conclusion that all attention manifests a sort of agency. The actionist’s claim to have discovered an essential link between attention and agency is thus on secure ground. But like a false belief or an unfulfilled intention, sometimes attention does not achieve what it aims at. In attending, we do not always avoid distraction. In such cases, interference within our perceptual and perceptuomotor states prevents us

from exercising perfect control over what we do and as a result action can suffer.

In the next chapter, I turn to address a traditional challenge to accounts, like my own, that have a certain reductive aspiration. In Chapter 1, I argued that attention is the process whereby an agent's intention settles the fine details of its own execution. In doing so, I suggested, the agent *acts*. In this chapter, I extended this claim to include not only intentions but the agent's overall priority system or total motivational state (warts and all). But what warrants the alleged reduction of an *agent's* control to control by a *motivational state*—whether an intention or the more encompassing total motivational state of which an intention is a fragment? How, in other words, can an account framed wholly in terms of mental states, events, and processes avoid depriving the agent of her rightful role in the generation of behaviour? I take up this question next.

Appendix

In this appendix, I want to sketch some implications of the priority view for how we should understand the functional profile of the “attentional” or “integrated priority map”. In the body of the chapter, I suggested that if the priority view is right, then the integrated priority map may play a less fundamental role in automatic attention than is sometimes supposed. My view is that the source state for automatic attention is the subject’s overall priority system, by which I mean total motivation. Further, my view is that the priority map itself instantiates a form of attention, insofar as it brings to bear the agent’s priorities on visual information to yield a priority ranking of visual information in the form of a map. I now want to expand on the implications of the priority view for how we should conceive of this map and its relation to emotions and affects. In doing so, I will assume that the priority map plays a mediating role in attention. I’ll approach these points by comparing how the environmental salience view and the priority view interpret the functional role of the map in, e.g., lateral intraparietal cortex in mediating the flow of attention.

Recall that whereas the salience map was said to be organized according to the distribution of salience in a scene, the priority map is organized according to the agent’s current overall priorities. If we understand the priority map to be a type of representation, then it is not a representation of the agent’s priorities, but of the relevance of locations in the scene to the agent given her priorities. As with the salience map, the priority map plays a *monitoring* or tracking role with respect to some property of the environment. And like the salience map was hypothesized to do, the priority map also *regulates* attention automatically in response to detection of this property—e.g., by biasing the processing of feature maps in visual ventral areas in ways that are sensitive to the map’s priority rankings. The main difference I’ve stressed between the salience map and the priority map as potential mechanisms of attentional control (or mediation) concerns the motivated character of the control. I’ve suggested that whereas the environmental salience view depicts automatic attentional shifts as a brute, reflexive response triggered by the detection of environmental salience, automatic attention shifts possess the profile of a motivated response on the priority view. On the priority view, the representation of a location as having relevance to

one's goals or priorities motivates turning one's attention toward it. Whereas a perception of environmental salience may be able to trigger selectively enhanced processing at a location, the representation of information as having high priority to one's goals motivates its prioritization.

By both monitoring incoming information for its relevance to the agent's goals and automatically regulating attention in light of its detected relevance, the integrated priority map, on this conception of it, bears a functional profile similar to that of a "valenced" representational state is sometimes said to possess.⁷ Examples of valenced states include affects like hunger and fatigue and emotions like fear and sadness. The nature of valent mental states is a matter of ongoing controversy, but, on a broad family of views, these states possess the dual role of automatically monitoring and regulating the presence of organism-environment relations pertaining to an organism's well-being ("concerns").¹ For example, fear serves to detect the relation of danger and, in virtue of its negative valence, to trigger responses aimed at eliminating its presence—e.g., by withdrawing, fleeing, etc. Pleasure in eating, by contrast, signifies that the attempt to replenish is going well and, in virtue of its positive valence, motivate the activity's continuation ("more of this!"). A valenced state, so understood, may simultaneously serve to monitor or detect relations of concern or need (e.g., danger, loss, hunger, success, etc.) and automatically to motivate responses that serve to regulate those relations (e.g., approach or avoidance). In a related vein, the integrated priority map could be understood as realizing a valenced role: to detect the relative priority of incoming visual information in relation to the agent's priorities and, in motivated response to these detections, automatically reorient attention.

The suggestion that the integrated priority map may be valenced or valence-like may go further than possession of a common structure. It might also provide an insight into instinctual or emotionally guided behaviour. Specifically, it is tempting to conjecture that the priority map and paradigmatically valenced states (e.g., emotions like fear and anxiety, affects like pain and pleasure, etc.) mutually enable each other in the following respect. The priority map, on its own, would have no role to play in the regulation of behaviour were there not a source of agential priorities feeding into it. The priority map depends on the agent's motivational priorities to serve as the basis for its priority rankings; otherwise, there would be nothing for visual information to be represented as having pri-

⁷The idea that valenced states function simultaneously to monitor and regulate relations of concern is a fairly standard claim in the literature on emotions (see, e.g., Prinz 2004; Cochrane Cochrane 2018). A terminological point: the term "concern" as used in this context does not refer to a mental state, but to the objective relations that the valenced mental state functions to monitor and regulate. Thus, danger, emotional loss, and physiological hunger are the "concerns" tracked and regulated by feelings of fear, sadness, and hunger, respectively. Alternatively, we might speak of the "needs" of safety, emotional connection, satiety, etc. In claiming that valenced states track such states, one does not thereby commit to regarding them as perceptual states.

ority *to* or *for*. If so, the priority map could not simply be one among other valenced states. Rather, to play its monitoring or detecting role, it requires input from motivation. At the same time, emotions and affects might have the opposite deficit. Although these states have no trouble monitoring for relations of concern (often doing so outside the focus of attention), they may need to harness the priority map if they are to regulate behaviour once a concern has been detected.⁸ By harnessing the priority map, a concern (i.e., one of the agent's priorities) becomes prioritized, and opportunities for its regulation (e.g., its removal, continuation, approach, withdrawal, etc.) become selectively enhanced. This points to a very basic role for attentional prioritization in enabling motivationally guided behaviour.

⁸Cf. Cochrane (2018): "Overall, my claim is that the most plausible output function of pleasant and unpleasant affect is to capture attentional priority" (55). Cochrane links this with the claim that affective states of pleasure and pain succeed in motivating behavioural response only by triggering attention to the pleasurable or painful sensation. However, I think Cochrane would reject my suggestion about the priority implementing a valence role.

Chapter 3

Attention and the Problem of the Disappearing Agent

3.1 Introduction

The next two chapters are about conscious attention. On the view that I'll defend, to attend consciously is to adopt a distinctive sort of point of view. I call this sort of point of view "motivational perspective". On my account, motivational perspective contributes both to our experience of ourselves as agents and to our perceptual experience of the world around us (including our own moving bodies). Its possession is consequently a genuine point of intersection between the phenomenology of agency and the phenomenology of perception. As I hope to show, failure to adequately grasp the character of this subjective perspective has the potential to distort our grasp of the relationship between these domains of experience. In the present chapter, I will be concerned with the role motivational perspective plays in the experiences of acting. In the next, I discuss how motivational perspective helps shape our perceptual awareness of the world.

In showing what difference conscious attention makes to the experience of acting, I am also showing what difference conscious attention makes to agency *itself*. In particular, I am concerned in this chapter with a traditional problem in philosophy of action called the "Absent" or "Disappearing Agent Problem". To see the problem, consider that an adequate theory of action must accommodate the intuitive difference between behaviours that you, as an agent, *do* and those that merely *befall* or *happen* to you. For example, it must account for the difference between the act of intentionally raising your arm and someone else raising your arm despite you. The Disappearing Agent Problem is the charge that a given theory or class of theories fails to accommodate that difference because it fails to accord the *agent* with the necessary causal role in acting. If this charge is correct, then a

theory that sets out to explain *doings* delivers only *passive happenings*.

In Chapter 1, I mounted a version of the above challenge to Wu's suggestion that attention realizes the role of the agent in action. Specifically, on Wu's account, attention plays the role of the agent in the selection of one from many alternate possibilities (see §1.6.2). Against Wu's account, I objected that it ultimately accords no role to the agent in action. It therefore confronts a Disappearing Agent Problem. At the time, I acknowledged that a similar problem might seem to arise for my alternative proposal that agential control is guidance by a motivational state of the agent. Specifically, one can ask: why believe that control by a *motivational state* amounts to control by *the agent*?

This chapter presents my answer to this question. I will argue that control by a motivational state can amount to control by the agent when the former plays the role of the agent's motivational perspective and that, furthermore, exercises of conscious attentional control realize such perspective. In doing so, I uncover an essential link between the motivational control of conscious attention, on the one hand, and an agent's ownership of action, on the other. I thereby establish fundamental explanatory connections between consciousness, attention, and agency.

However, there is an initial barrier to pursuing these connections. Beyond rough and ready glosses like the one offered a moment ago, philosophers use the expression "Disappearing Agent Problem" and its cognates to describe distinct explanatory challenges. And, as we'll see shortly, not all these challenges are equally compelling. Among those versions that mount a genuine explanatory challenge, some pose a more fundamental challenge than others. Accordingly, my first task (in §3.2) is to isolate the most fundamental, genuine, and outstanding explanatory challenge to accounts, like mine, that claim to reduce agent-causation to a species of motivational state-causation. This challenge, I suggest, is to accommodate a central phenomenal property of the experience of acting: the sense of first-person ownership of action. Having identified the Disappearing Agent Problem in its most fundamental form, the remaining sections develop and defend my positive account of this phenomenal feature of agentic experience. I do so first for the case of "immersive" experiences of acting (§§3.3–3.4), and then extend it to non-immersive cases (§3.5). In closing (§3.6), I explain how my approach hopes to accommodate various pathological or nonstandard experiences of acting (through a comparison with opposing "representationalist" accounts of these experiences) and draw broader morals regarding the role of consciousness in being an agent.

3.2 “The Disappearing Agent Problem”

According to “reductive” or “event-causal” theories of action, we can exhaustively characterize action in terms of a specified class of mental events (understood liberally to include states and processes), behaviours, and causal relations between them. Typically, the relevant class of psychological events are those that motivate and rationalize the agent’s behaviour. For example, on Davidson’s theory, an action is an event that is intentional, and an event is intentional iff caused by reasons that motivate and rationalize it—e.g., a desire to Φ and a belief that Ψ -ing is a way to Φ (Davidson 1963). Whereas Davidson’s theory identifies the action with the *product* of a certain process of mental causation (i.e., the resultant behavioural event), other reductive accounts identify action with the *process* consisting of the mental event *causing* the behavioural event (e.g., Dretske 1988, Ch.2; Searle 1983, Ch.3). Each type of account qualifies as “reductive” or “event-causal” insofar as it explains the actions of an agent wholly in terms of the causation of bodily movement by mental states or events of the agent (e.g., her intention).

But why should we believe that the causation of bodily movement by mental states of the agent amounts to *the agent* moving her body, i.e., acting? This is the question that animates the Disappearing Agent Problem against reductive, event-causal theories of action. As Velleman (1992) summarizes the worry, “the standard story of action” (approximately, Davidson’s) “fails to include an agent or, more precisely, fails to cast the agent in his proper role. In this story, reasons cause an intention and an intention causes bodily movement, but nobody—that is, no person—does anything” (461). Such a view could at most describe a “defective” sort of agency in which:

... various roles that are actually played by the agent himself in the history of a full-blooded action are not played by anything in the story or are played by psychological elements whose participation is not equivalent to his. In a full-blooded action, an intention is formed by the agent himself, not by his reasons for acting. Reasons affect his intention by influencing him to form it, but they thus affect his intention by affecting him first. And the agent then moves his limbs in execution of his intention; his intention doesn’t move his limbs by itself. The agent thus has at least two roles to play: he forms an intention under the influence of reasons for acting, and he produces behaviour pursuant to that intention. (Velleman 1992, 462)

The difficulty with the standard Davidsonian story, Velleman suggests, is that none of the psychological elements that it posits are suitable to occupy the functional roles that we intuitively attribute to the agent in relation to his reasons, intention, and behaviour—e.g., of *considering* reasons, of *forming* an intention to act on the basis of considered reasons, and

finally of *executing* that intention in pursuit of his intended goal. On such a view, the agent seems to *disappear* from the story altogether. What is left are passive causal connections between bodily movements and various mental events that the agent undergoes. This, briefly and roughly, is the “Disappearing Agent Problem” for the standard Davidsonian account of intentional agency.¹

Beyond this relatively neutral formulation, philosophers elaborate the worry in markedly different ways. As noted earlier, in my view some elaborations are more compelling than others. I will begin (in §3.2.1) by briefly discussing a version of the Disappearing Agent Problem that does not, to my mind, present a genuine *objection* to reductive accounts. I then turn to discuss a version that does pose a genuine challenge but that, in my view, is not the most fundamental challenge in the neighbourhood. This will pave the way for what I take to be the deepest challenge in the vicinity.

3.2.1 Ambitious formulations

According to the first version of the Disappearing Agent Problem I will consider, there are *a priori* reasons why action—and so the relation of ownership that agents stand in to their actions—cannot be analyzed in purely event-causal terms. As some authors make the point, event-causal accounts like Davidson’s invariably identify the agent with an entity of the wrong metaphysical category—e.g., a mental state or event.² A related complaint is that event-causal accounts inevitably replace the agent’s *acting* or *doing* with a mere

¹It is important to distinguish the Disappearing Agent Problem from a different problem sometimes raised against reductive theorists: that of deviant causation (Davidson 1973). Those mounting the Disappearing Agent Problem to event-causal theories tend to regard the former as the more fundamental of the two problems. For example, Velleman (1992) writes:

[T]he discussion of “deviant” causal chains has diverted attention from simpler counterexamples, which omit the agent without lapsing into causal deviance; and it has thereby engendered a false sense of confidence in the requirement of causal normality, as sufficient to protect the standard story from counterexamples. In reality, an agent can fail to participate in his behaviour even when it results from his motives in the normal way. Consequently, no definition of causal normality will fix what ails the standard story. (1992, 464)

²This appears to be one of Michael Brent’s (2017) central objections to event-causalism about action in a recent defense of non-reductive agent-causalism:

[I]f causation is a relation among discrete events, and if your mental events are causing the relevant movements of your body during an action, then, in order to ensure that you are playing a causal role when moving your body during an action, proponents of typical [event-causal] solutions must claim that you are *identical with* the causally relevant mental events. However, it is implausible on metaphysical grounds that you are identical with your mental events, so it follows that your mental events are not causing the relevant movements of your body during an action. Thus, typical solutions to the problem of action are in trouble (Brent 2017, 660-1, italics original).

“happening” in relation to which the agent is merely a passive “subject”, “witness”, or “bystander”, or that they replace the agent with a mere “arena” within which internal events interact like billiard balls to cause behaviour.³ These versions of the Disappearing Agent Problem make the ambitious claim that *any* attempt to reduce the action of an agent to a species of event-causation necessarily results in failure to describe *action* at all. It at most describes one event passively causing another. For this reason, I call these “ambitious formulations”. Perhaps unsurprisingly, those endorsing ambitious versions of the Disappearing Agent Problem typically espouse some form of primitivism about agent-causation, with the latter often (though not always) being explained in terms of the notion of substance-causation (see e.g., Alvarez and Hyman 1998; Hornsby 2004; Steward 2012; Brent 2017).

If ambitious formulations were successful, they would pose a very fundamental problem for reductive theories of action. However, ambitious formulations do not seem to present a genuine challenge. This is because ambitious formulations either mischaracterize the event-causalists’ position or are unduly dismissive of it. Consider the claim that reductive proposals wrongly imply that the agent is an event. Event-causalists can reply that their position lacks the putative implication. They can even accept that the agent is an object or substance of some kind and that, in acting, this substance moves itself or its body—i.e., causes itself or its body to move. They will simply add that when an agent acts, she does so in virtue of certain mental states or mental properties (cf. Mele 2003, 225). This is analogous to the claim that when a rock shatters a window, it does so in virtue of certain of its causally relevant properties (e.g., its mass and velocity). The objection that event-causalists misclassify the agent as an event therefore misunderstands their position. On closer examination, the second variant of the ambitious challenge fares no better than the first. In response to the claim that event-causalists invariably replace actions with “mere happenings”, the reductionist can reply that such passive language is question-begging (cf. Schlosser 2011, 23). An individual’s involvement in an event or happening is not clearly sufficient to render the individual a passive “victim” to the event. Nor is it clearly sufficient for the event to “befall” or “overwhelm” the individual in question. Indeed, event-causalists have proposed *criteria* for distinguishing happen-

³Melden (1961) writes: “It is futile to attempt to explain conduct through the causal efficacy of desires—all *that* can explain is further happenings, not actions performed by agents” (128-9, italics original). When reflecting on the world as an event-causal order, Nagel (1986) writes that “*my doing* of an act ... seems to disappear when we think of the world objectively. There seems no room for agency in [such] a world... there is only what happens” (110-11, italics original). Hornsby (2004) writes “I have claimed that there is alienation of an unthinkable sort when an agent is portrayed as if she were *merely an arena for events*” as Hornsby claims the agent is portrayed on all event-causal accounts (185). Hornsby also writes that “human agents are not merely *things within which things happen*, and they clearly *do* play a role in the arena within which their actions are found” (176, italics added). Many more examples could be provided.

ings that are intuitively passive (e.g., reflexes and externally generated movements) from those that are exercises of an agent's control. To claim that no attempt to accommodate the pre-theoretical active-passive distinction in event-causal terms can succeed is not to engage seriously with the event-causalist's position.

For these reasons, we should be wary of ambitious formulations of the Disappearing Agent Problem. While there may be *other* grounds for rejecting event-causal accounts, ambitious versions of the Disappearing Agent Problem seem not to provide a genuine objection. Accordingly, I set aside ambitious versions of the challenge.

3.2.2 Velleman's Problem (the Shrinking Agent Problem)

Like ambitious theorists, Velleman (1992) argues that a Disappearing Agent Problem confronts the standard Davidsonian story of action. Unlike the ambitious theorists, though, he does not object to the event-causalist's reductive aspiration. Velleman ultimately shares this aspiration. Rather, Velleman claims that the problem is with the proposed reduction. Furthermore, as Velleman proceeds to explicate the difficulty, the Disappearing Agent Problem for the Davidsonian story is not that it fails to explain action of *any* sort. Rather, in Velleman's view, the problem with the Davidsonian story is that it can explain only specific "defective" cases of action, such as the motivated activities of mere wanton addicts.⁴ "Wantons" are agents who engage in motivated activity, but who never—whether through lack of concern or lack of ability—consciously reflect on the quality or merits of their motives. They can be subject to competing motives—e.g., a desire to take a drug and a simultaneous desire to not take the drug—but the outcome of such competitions is entirely settled by the relative *strength* of the competing motives.

As Velleman sees things, the Davidsonian story is adequate to describe the motivated activities of unreflective wantons. But he thinks Davidson's story fails to explain the actions of agents who attempt to *intervene* upon motivational competitions in order to bias the outcome in a particular direction. Following Frankfurt, Velleman suggests that a hallmark of distinctively "human" agents (or "persons" in Frankfurt's terminology) is that, when faced with opponent motivational tendencies (e.g., between the desire to take the drug now and the desire to quit), the agent is typically not a passive bystander or witness to the competition occurring inside him (as the wanton is). Rather, the agent tends to "put his weight behind" the motives that he "endorses" or "identifies with" and to

⁴This is at least how Velleman explains the notion of defective or lower-order agency in Velleman (2000b). In doing so, Velleman appeals to distinctions familiar from the work of Frankfurt (e.g., 1971; 1976), including Frankfurt's trichotomy between actions, motivated activities, and passive bodily movements. In Velleman's terms, Davidson's account can at most accommodate Frankfurt's intermediate category of motivated activities.

suppress the influence of those he disavows as “external” to himself (cf. Frankfurt 1971). Through processes of reflective identification (and dis-identification), Velleman and others in this tradition suggest, it becomes possible for the objectively *weaker* motive (e.g., to persevere in one’s resolution to quit) to triumph over the objectively *stronger* one (e.g., to take the drug now). Where an agent’s act of identification with a given motive is successful in generating behaviour congruent with the endorsed motive, the agent exercises self-control (or “autonomy”). This self-control is said to yield a form of psychological *ownership* over one’s own conduct that mere wantons would never experience, given the latter’s lack of concern over the quality of their own motives. Relatedly, where the agent’s act of identification *fails* to produce behaviour congruent with the identified-with motive, the agent is thereby *alienated* from his conduct. In such cases, the agent experiences his behaviour as not fully his own. As Frankfurt puts it when describing someone whose drug addiction has overpowered his decision to quit the drug, such an unwilling addict “finds himself a helpless or a passive bystander to the forces that move him” (Frankfurt 1971, 17).

Finally, the Disappearing Agent Problem, on Velleman’s conception, is to explain the crucial psychological processes of reflective identification and disavowal without illegitimately *presupposing the agent* as the active director or causal force behind these psychological processes. If such a reductive account could be given, Velleman supposes, then the event-causalist will have demonstrated how it is possible for agents not merely to be *caused* or *moved* to act by states to which they are passively subject (in the manner of wanton addicts). The event-causalist will have further shown how it is possible for agents to exercise genuine *self-control* and so to be truly active in the generation of behaviour. Because Velleman’s objection to the Davidsonian does not concern the possibility of agency *per se*, but a specific subset of human agential capacities, Mele (2003, 220) proposes renaming Velleman’s puzzle the “Shrinking Agent Problem”.⁵

Velleman is right that self-control presents a genuine explanatory challenge to the original Davidsonian story. Event-causalists must explain not only the activities of unreflective wantons, but the autonomous acts of persons. However, there is also a difficulty with formulating the Disappearing Agent Problem as Velleman does. The difficulty is that it exaggerates the successes of the Davidsonian story from the very start. Velleman assumes, wrongly, that the only situation in which the standard Davidsonian story requires amendment in order to accommodate an agent’s “ownership” of her actions is one

⁵Velleman proceeds to offer a reductive solution to his puzzle in terms of a higher order motive to “make sense” or conform to one’s understanding of oneself. See also Mele (2003, Ch. 10) for a reductive proposal invoking a higher order motive to act on the basis of one’s strongest reason. On Mele’s view, this motive, when instantiated, enters into the “positive motivational base” of the endorsed first-order motive, increasing its motivational strength (222).

in which the agent's behaviour is controlled by a self-consciously endorsed motive (and is controlled partly because of the agent's reflective endorsement of the motive). Similarly, he assumes that the only type of *disownership* in need of explanation is the sort exemplified when behaviour is guided by a motive that the agent self-consciously rejects as not part of herself. Consequently, Velleman concludes that if we can simply enrich the standard Davidsonian story with an account of self-control and its potential losses (in, e.g., unwilling addiction), we will have remedied the deficiencies of the standard story. Against this, I will now argue that the Disappearing Agent Problem is in fact more fundamental than Velleman acknowledges. This is because the problem of explaining agential ownership arises earlier. In particular, it arises for even rudimentary manifestations of agency, including those of unreflective wantons.

We can motivate the problem by way of a contrast between two sorts of "alienation" experiences (cf. Schlosser 2011, 25-6). On the one hand, there is the type of alienation that unwilling addicts experience when they are moved to act on the basis of motives they reflectively disavow. On the other hand, there is the more radical type of alienation experience reported by patients suffering from, for example, anarchic hand syndrome and schizophrenic delusions of control. In the empirical literature, the latter are often called "passivity experiences" (Spence et al. 1997; Blakemore et al. 2000). For example, after performing a task involving intentional finger movements, a schizophrenic patient suffering from experiences of alien control reported: "I felt like an automaton, guided by a female spirit who had entered me during [the movement]" (Spence et al. 1997, 2001). While Frankfurt, Velleman, and others may be right to regard unwilling addicts as in *some* sense passive in relation to their addictive behaviour, it's equally clear that unwilling addicts *retain* a sense of agency over their addictive behaviours: a sense of agency that the patient experiencing alien control reports lacking. Presumably even unwilling addicts experience exercising conscious control as they deliberate about, for example, how they will obtain their next hit, even if they don't self-consciously endorse or approve of their intention to take the drug. No less than their willing counterparts, unwilling addicts experience themselves as *considering* reasons for and against different courses of action (e.g., about whether to look for the drug in location A or in location B), *forming intentions* on the basis of earlier deliberations (e.g., to look for it at location A), and finally *executing* these intentions by moving their bodies. Any account of unwilling addiction that is inconsistent with these points would be guilty of attributing an excessively alienated form of agency to the unwilling addict—one more akin to the passivity experiences of alien control and anarchic hand patients.⁶

⁶Cf. Hornsby (2004), who, in response to Velleman, writes:

But a person who appreciates that her conduct is out of accord with what she values, or is

This suggests that there is a sense of action ownership that is not only different, but psychologically more basic than the sense of reflective ownership that distinguishes self-controlled actions from less self-controlled ones. The former variety of ownership experience is more basic than the latter because it does not depend on metacognitive abilities, such as the ability to consciously reflect upon one's own motivational states, to evaluate whether they provide one with strong reasons for action, and finally to endorse or disavow them on the basis of one's assessment of their relative merits. Consequently, the sense of ownership that is retained in unwilling addiction will likely also be enjoyed by mere wantons. A wanton's states of captivation by her first-order motives, and the accompanying immersion and engrossment with which she pursues their fulfillment, may disqualify her from the kind of ownership and alienation experiences that self-controlled and unwilling addicts experience. But it presents no barrier to experiencing a more minimal variety of action ownership that is present automatically in action rather than actively conferred on action through endorsement of it as "truly one's own" (cf. Schlosser 2011, 26). For these reasons, I call this more basic form of ownership experience "basic subjective ownership for action" or "basic subjective action ownership".

The requirement to accommodate basic subjective action ownership is, I suggest, the most fundamental, genuine Disappearing Agent Problem facing extant event-causal accounts. For the remainder of this section, I consider two potential responses one might offer to my claim to have identified a *genuine* challenge to extant reductive accounts.

3.2.3 Already owned?

The first response to my claim to have identified a genuine challenge for the event-causalists accepts that event-causalists are required to accommodate basic subjective action ownership, but argues that extant forms of event-causalism already meet this requirement. Specifically, one might argue that agents experience their actions as their own when there is a rationalizing causal explanation of those movements in terms of the agent's

swayed by factors whose influence she regrets, admits her own motivations even if she does not approve of them. The desires and emotional states which explain what she does are after all states of *hers*—of the human being whose capacities to make movements are exercised—and, even where she feels alienated from them, they are not adventitious forces in her brain. (To think of adventitious forces in the brain seems more appropriate in understanding, say, the involuntary movements of sufferers from anarchic hand syndrome, which lack any personal psychological explanation) (Hornsby 2004, 182-183, italics original).

Hornsby raises an important objection here to Velleman. However, she appeals to it in the course of mounting a more ambitious Disappearing Agent Problem against Velleman's account, similar to the kind discussed in the previous section. In my view, her observation instead suggests that there is a more basic variety of ownership experience which Velleman has incorrectly assumed to already be adequately understood on the Davidsonian picture.

conscious motivational states—e.g., conscious intention. Consequently, where agents experience a movement as *not* their own doing, this is because there is no rationalizing explanation for the behaviour in terms of the agent’s conscious attitudes.⁷

To elaborate this suggestion, we might appeal to Searle’s (1983) thesis that conscious intentional actions are accompanied by an “experience of acting” and moreover that this experience of acting is identical with an “intention in action” (87). As he explains, an intention in action is the phenomenally conscious, present-directed mental state responsible for initiating and guiding an agent’s behaviour when she acts intentionally. Searle further holds that an intention in action possesses token-reflexive intentional content along the lines of (very roughly) <my arm rises as a result of this intention in action> and the content of this state is satisfied iff the intention in action brings about the behaviour it represents (and does so “continuously” and with “plannable regularity”) (see Searle 1983, 136 ff.). Whereas the canonical means of linguistically expressing a “prior” (or “future-directed” or “distal”) intention is “I will do A” or “I am going to do A”, Searle suggests that the canonical mode of linguistic expression for an intention in action is “I am doing A” (Searle 1983, 84). Drawing on this suggestion, one might propose that wherever bodily movements are controlled by one’s intention in action to move one’s body, one experiences those movements as one’s own doing. And where one experiences disownership toward one’s movements, the movements are controlled by something other than an intention in action. Significantly, on this view, the same psychological elements that account for the motivated and rationalized status of the act simultaneously secure basic subjective ownership for the action—namely, the intention in action.

Searle’s proposal seems to generate the correct prediction for some passivity experiences. For example, in anarchic hand syndrome, subjects report that their bodily movements (e.g., grasping a piece of food from another person’s plate) are incongruent with their consciously expressed goals or intentions (e.g., to eat only from their own plate). Often, this is reflected in the subject actively attempting to suppress the hand’s anarchic movement (e.g., by sitting on it). In such cases, one might plausibly infer that if there is an intention or motor command guiding the subject’s movements at all, this state is not an intention in action in Searle’s sense. This is because it is introspectively inaccessible to

⁷While Hornsby would not endorse the strategy I am now sketching on behalf of the event-causalist (because she rejects event-causalism on ambitious grounds), we might recall her claim that “To think of adventitious forces in the brain seems more appropriate in understanding, say, the involuntary movements of sufferers from anarchic hand syndrome, which lack any personal psychological explanation” (Hornsby 2004, 183). Hornsby hints here that what explains the *alienated* status of the anarchic hand patient’s behaviour (as opposed to the *disavowed* status of the unwilling addict’s actions) is the absence of “any personal level psychological explanation” for the movements (e.g., in terms of a conscious intention to move her arm). This is the suggestion I am now exploring on behalf of the event-causalist in accommodating basic subjective ownership.

the subject as well as functionally unintegrated with the subject's other attitudes. If so, then Searle's account is consistent with the finding that anarchic hand patients experience their behaviours passively. The experientially alienated status of these behaviours reflects their unconscious and subpersonal source.

However, Searle's account does not fare as well in other cases. Consider that patients experiencing schizophrenic delusions of alien control are often able to follow task instructions successfully. They also recognize that their behaviour is congruent with their explicit goals and show no tendency to desist or intervene upon behaviour once it is under way.⁸ On the face of things, this provides evidence that, unlike in anarchic hand patients, it is the schizophrenic patient's *conscious* intention to execute task instructions that controls—i.e., initiates, sustains, and guides—behaviour. And yet these patients report feeling disownership over their actions. They report experiencing their actions as possessing an outside source. Against Searle's proposal, these agents appear to feel no ownership of their actions despite these actions being under the control of conscious intentions.⁹

The evidence from schizophrenia suggests that it may be possible for a conscious intention to control an agent's action without the agent thereby feeling ownership of the action. Some have attempted to defend related conclusions on purely introspective grounds. In particular, some have challenged the assumption that even the *experience* of motivational causation is ever sufficient for experiencing oneself as acting. For example, Horgan and colleagues (Horgan et al. 2003; Horgan 2015) ask us to imagine scenarios in which one is introspectively aware of a motivational state to behave somehow—e.g., of a conscious urge to clench one's fist—and one is also aware of this motivational state *causing* the motivated behaviour—e.g., a fist-clenching. Under the described conditions, they suggest, one would not experience oneself as *acting* or having acted. In particular, insofar as the felt source of one's behaviour is the *felt urge* rather than *oneself*, Horgan and colleagues suggest that one's experience of the motivated behaviour—e.g., the urged fist-

⁸As Frith et al. (2000) note, for example, one of the distinguishing features of schizophrenic patients (in marked contrast with anarchic hand patients) is that "In most cases the actions made when the patient 'feels' that he is being controlled by alien forces are not discrepant with his intentions. Thus the patient may be correctly performing the task set by the experimenter (e.g. making random movements of a joystick) at the same time as having the experience of passivity (Spence et al. 1997). The patient does not try to correct these 'controlled' actions or prevent them from occurring." (Frith et al. 2000, 1784).

⁹One might object that while the *prior* or *distal* intentions (e.g., to execute a certain task in the near future) of these patients are conscious, it is not obvious that their *present-directed* intentions in action are conscious. The point is a fair one, though the cases discussed in the next paragraph sidesteps this particular worry. Furthermore, if we examine the patient's subjective report quoted earlier, we see that the subject feels as though the alien agent control began *during* the movement, rather than at its outset. If the agent's present-directed intention were unconscious, one might expect the subject to start experiencing alien control at the beginning of the movement (which is when the present-directed intention commences). Rather, the passivity experience appears to commence part way through.

clenching—will remain subjectively *passive* (as well as “strange and alienating”) (Horgan 2015, 35). It would be, they suggest, akin to watching an external causal process unfolding in nature. While their specific example is an urge, they take the observation to generalize to all motivational states, including presumably intentions. They conclude that an experience of acting does not “present your own behavior to you as having your own occurrent mental events as its source” (Horgan et al. 2003, 225). Instead, when you act, the apparent source of your behaviour is *you, the agent*. They call this the “phenomenology of self-as-source” and claim that it is a key phenomenal component in ordinary, non-pathological experiences of acting.

If Horgan and colleagues are right, then it seems Searle crucially mischaracterizes the experience of acting when he characterizes its intentional content as, for example, <*this intention-in-action to raise my arm causes my arm to rise*>. To experience one’s own behaviour as controlled by one’s own mental states—even one’s own intentions—does not amount to experiencing *oneself* as controlling one’s behaviour. Indeed, Searle himself seems implicitly to acknowledge as much when he characterizes the canonical linguistic expression of an intention in action as “*I am doing A*”, rather than the token-reflexive content that he in fact attributes to the intention in action. Consequently, I think we can plausibly reject the suggestion that basic subjective action ownership exists wherever the action is the product of control by a personal level motivational state (or the process of the motivational state causing motivated behaviour).¹⁰

In §3.6, I’ll return to examples of alienated experiences of acting, both hypothetical and actual, and propose an alternative explanation of them. For now, I wish to draw a tentative moral from the possibility of phenomenologically passive motivational state-causation. First, extant event-causalist accounts appealing only to control by the agent’s attitudes (e.g., intentions) fail to preclude profoundly alienated experiences of acting, both in actual clinical populations like schizophrenia and in purely hypothetical examples like Horgan et al.’s. They are consequently vulnerable to the objection that their view falsifies the experience of acting. Second, we can identify the principal feature of extant event-causalist proposals that leave them open to this challenge. Specifically, there is nothing to preclude a conscious intention from controlling the agent’s behaviour while the agent experiences herself as observing the control process “from the outside”. A more adequate event-causalist account of basic ownership would therefore invoke mental elements whose deployment within the control of behaviour forecloses the possibility of such subjective detachment. These would be mental elements whose deployment in behavioural

¹⁰For further criticisms of Searle’s intention-based account of the experience of acting, see Bayne and Levy (2006); Bayne (2008, 2011); Mylopoulos (2017) and Shepherd (2017a). Kriegel (2015) defends an account of the experience of acting much like Searle’s, which he frames in terms of “phenomenal tryings”. It is vulnerable to many of the same objections as Searle’s.

control are sufficient for the agent to experience herself as immediately engaged *within* the process of behavioural control as an active participant rather than as a spectator observing the process as it unfolds from *without*. In later sections, I put forward an account that meets this explanatory demand and thereby solves the Disappearing Agent Problem in its most fundamental form.

3.2.4 Yes, this is agency (not *merely* sense of agency)

Let's turn to the second response one might give to my claim to have identified a genuine explanatory challenge to event-causalists. This response grants that extant versions of event-causalism have not accounted for the subjective experience of basic action ownership. However, it then argues that they are under no burden to accommodate such an experience. Whereas basic subjective action ownership is a matter of the *sense* of agency, event-causal accounts are attempting to provide an account of agency *itself*. And it is no objection to an account if it "leaves out" what it has never claimed to explain. Even if an account like Davidson's, for example, fails to account for the phenomenology of acting, it might still be fully adequate as an account of action itself.

The above response in effect claims that an account of the sense of agency isn't a *requirement* on an account for action. A similar intuition will probably incline one toward affirming the converse point: namely, that an account of the sense of ownership is not *sufficient* for an account of action ownership. Just as we will not have explained the nature of a material object (e.g., an apple) by giving an account of our representation of it (e.g., how it visually appears to us), similarly we will not come closer to understanding the relation of ownership that agents bear to actions simply by investigating their *experience* of that ownership relation. An account of ownership should tell us *what it is* for an agent to act, which means giving an account of the *fact* of action ownership. And one cannot meet that explanatory demand by providing an account of how it subjectively *feels* or *appears* to one to act—i.e., to give an account of the *sense* of action ownership.

In response, this would be a decisive objection to the approach that I am taking in this chapter if it were true that questions about action ownership can be entirely separated from questions about the phenomenology of action ownership. It would be decisive to my approach because I am indeed proposing to account for the metaphysical relation of action ownership by attempting to accommodate its phenomenology. The reader will accordingly have grasped my account of subjective action ownership and its motivations over rival accounts once they have grasped why we *cannot* treat the nature of action ownership as wholly independent of its phenomenology. As we shall see, the relation between the fact of ownership and the sense of ownership is much more intimate than the above

response depicts it as being.

Let's take stock. I have identified the most fundamental, genuine version of the Disappearing Agent Problem confronting extant event-causal theories of action. This is to accommodate basic subjective action ownership. I will now answer this challenge. I will do so within the context of my thesis from Chapter 1 that agential control is guidance control by a motivational state of the agent. I will defend my account of basic subjective ownership on the basis of its ability to solve a puzzle that arises as soon as we begin to reflect on how the agent manifests subjectively within "immersive" experiences of acting—i.e., experiences of acting that even an unreflective wanton is apt to enjoy.¹¹ After motivating my attentional account of basic subjective ownership as providing a solution to this puzzle, I elucidate the proposal by reference to the structure of attention as a form of conscious control.

3.3 A puzzle about the sense of ownership in "immersive" action

In this section, I present a puzzle about how the agent figures subjectively within immersive experiences of acting. After considering some unsuccessful solutions, I sketch my own solution in terms of conscious attention. I begin by drawing a couple of distinctions.

To begin, we must draw on a distinction between two types of ownership experience. These experiences are distinguished by what the subject feels ownership *toward*. First, one sometimes experiences being the *subject* of a bodily change. This experience is salient in experiences of passive or involuntary movement, as when one shivers from the cold, trembles in fear, or is tickled. If these movements are consciously perceived at all, one will usually experience them as changes in oneself (or one's own body). Second, one sometimes experiences being the *agent* of a bodily change. Usually when one acts—e.g., lifts one's arm—one's sense of ownership extends beyond the bodily changes that one undergoes to include the process of causing and controlling those bodily changes. In these cases, one not only experiences one's own body moving. One also experiences oneself *making it move*. Put differently, one experiences ownership of the act of moving one's body. The latter is a "first-person experience of acting".

Some experiences of acting are "immersive" in the sense of being characterized by heightened sense of engagement, absorption, or engrossment in a task or activity.¹² Plau-

¹¹See Velleman (2008) for the relationship between the wanton and the sorts of immersion experiences I will discuss in §3.3.

¹²I draw the terminology of "immersed" versus "detached" experiences of acting from Marcel (2003) and Pacherie (2008).

sibly, immersion experiences can differ along several dimensions of variation, including in their affective tone (e.g., whether boring or arousing, neutral or pleasant, etc.) and volitional character (e.g., whether characterized by subjective exertion or relative ease). The most well-known examples of immersion experience are so-called “flow” states (Csikszentmihalyi 1990; Csikszentmihalyi et al. 2005). Agents acting in flow frequently report (retrospectively) “losing themselves” in an engrossing activity, including a dance or a musical performance. We instinctively associate these immersive experiences with the highest achievements of skilled agency and flexible control—e.g., with a record-breaking athletic performance or an especially creative and dynamic play.

Immersive experiences of acting, like those enjoyed in flow, raise a puzzle about the sense of agency. The puzzle concerns how the agent, as such, manifests in the experience of acting. On the one hand, for an immersion experience to be an experience of *acting*, it seems that the agent’s behaviour must have its apparent causal source in the agent. Otherwise, the experience would not be an experience of acting at all, but merely of one’s own body moving passively. If it is to be an experience of acting, then, the agent must be subjectively *present* as the source of the movement. On the other hand, what subjectively distinguishes an immersed experience of acting from a more detached one—i.e., one in which one mentally “steps back” in order to monitor oneself or one’s performance—is that, in immersion experiences, agents feel they lose themselves in the act. To feel genuinely immersed rather than detached, it seems the agent must be *absent* from the experience. The puzzle is to explain the agent’s simultaneous presence and absence from the same experience.

Some will be tempted to solve this puzzle by rejecting the first of the inconsistent-seeming claims: that immersive flow experiences are genuine experiences of acting. As evidence, they might cite subjective reports like this rock-climber’s: “when things become automatic, it’s like an egoless thing, in a way. Somehow the right thing is done without you ever thinking about it or *doing anything at all ... It just happens*” (Csikszentmihalyi 1990, 62-3, italics mine). If agents in flow claim not to experience themselves “doing anything at all”, one might doubt whether this is really an experience of *acting*.

An initial reason to be dissatisfied with this response is that agents in flow typically report feeling *heightened control* over their behaviour and environment. One dancer reports: “I feel enormous power to effect something of grace and beauty”, and a chess player says: “I have a general feeling ... that I am in complete control of my world” (ibid. 59-60). Indeed, this feeling of control is thought to explain why flow-engendering activities are enjoyable to the point of being addictive. It isn’t clear how to interpret such reports except as expressing a feeling of *oneself as the agent* of heightened control. If so, then acting in immersion is, in at least in some instances, characterized by sense of agency.

Against this, one might try to accommodate subjective reports of heightened control by distinguishing the experience of control *possession* from the experience of control *exercise*. Specifically, one might suggest that while an agent acting in flow feels *in control* over what is happening with her body, she lacks any experience of *exercising* control. Provided experiences of acting entail experiences of exercising control rather than merely possessing control, flow experiences aren't genuine experiences of acting but perhaps of being poised or ready to act if necessary.

In response, it's true that agents in flow often emphasize control possession in their subjective reports. However, there's a crucial respect in which the same subjects report experiencing themselves as exercising greater control than they usually do. In particular, these subjects report feeling highly *focused*. For example, the rock-climber that I quoted earlier—the same one who said that, in flow, her behaviour “just happens” without her “doing anything at all”—finishes her statement by saying “And yet you're more concentrated”. With that caveat, the climber acknowledges that, even in flow, she is aware not merely of being able to control what she's doing, but also of allocating a certain kind of control to her performance. In particular, she experiences herself directing her performance by focusing her attention on the task at hand. This experience of heightened concentration on an activity seems to suggest that flow experiences are genuine experiences of acting.

If flow experiences are experiences of acting, one might wonder why subjects often describe them in intuitively passive terms. Let's set aside the possibility that the representational content of the experience is either paradoxical—e.g., in the way the waterfall illusion is sometimes suggested to be (Crane 1988)—or alternates between experiences of activity and passivity. Neither of these possibilities seems plausible on its face. The answer may come from the fact that experiences of acting possess multiple dimensions of variation, with flow occupying an unusual location in this space. Typically, when one performs at the upper threshold of one's skill level (not to mention in a potentially dangerous situation, as in some flow-inducing activities) one will have experiences variously of: self-evaluation about the quality of one's performance, anxiety about failure, and, perhaps most of all, effort or volition. Insofar as flow experiences lack many or all of these typical marks of agentive experience, they are apt to be reported as seeming agentively sparse. Nevertheless, as noted, there is one dimension along which these subjects experience directing more control to their task than they would otherwise: they experience greater concentration or focus. Partly for this reason, the experience is sometimes called “relaxed concentration”: a state of being calmly, effortlessly, unreflectively, yet attentively responsive to a challenging task situation.¹³

¹³For more on the difference between attention and effort (and their history of being conflated), see Bruya

Supposing, then, that immersion experiences are genuine experiences of acting, we need a solution to our puzzle that can accommodate the agent's felt presence as the source of control. One natural suggestion at this point would be to adopt a representationalist account of the experience of acting. According to this type of account, the phenomenal character of an experience of acting can be characterized, at least in part, by how the experience represents the world as being—in particular, by how the experience represents what the agent is currently doing (Peacocke; 2003; Bayne and Levy; 2006; Bayne; 2011; Mylopoulos; 2017; Shepherd; 2017a).¹⁴ To accommodate the agent's felt presence in the experience of acting, one might hold that the experience explicitly represents the agent via some indexical component. Mylopoulos endorses a version of this representationalist strategy when she writes:

The deployment of the essential indexical in their content makes it so that agentive thoughts explicitly represent oneself, as such, as the agent of the action in question. More specifically, if I have the thought that I am Φ -ing, the essential indexical makes it the case that my thought represents the thinker of that very thought as Φ -ing. In this way, an intimate link is forged between the subject of agentive awareness and the agent of the action, which might very well yield a robust sense of oneself as the source of that action. (2017, 557)¹⁵

Mylopoulos's idea is that we should explain subjective action ownership through the explicit representation of the agent as such within the representational content of a judgment. Agentive experience is revealed as a mode of explicit self-awareness whose content we might express by saying "I (the subject of this experience) am the agent of this action".

A natural worry to raise for the above strategy is that by requiring that agentive experiences explicitly represent the agent through the deployment of a first-person indexical, it cannot accommodate the *second* of our puzzle-generating claims about immersion

and Tang (2018) and Bruya (2010).

¹⁴An internal debate among representationalists concerns the format of the representation they take to realize agentive experience. Bayne (2011) argues that the representation is sensory and low level, whereas Mylopoulos (2017) argues that the representation is a belief or judgment that one is acting a certain way (see also Peacocke 2003). Shepherd (2017a) develops a distinctive view of the relevant representational state as "multi-categorical"—a product of the representational integration of the subject's conscious intention to move her body and of a concordant perceptual experience of her body moving. (Note: Shepherd doesn't explicitly advertise his account as a version of representationalism about experiences of acting. But that is nevertheless what his view appears to be. As I read Shepherd's proposal, the phenomenology of an experience of acting consists in a representation of one and the same event as being jointly a trying-by-me and a movement).

¹⁵Similarly, Peacocke writes:

The awareness of the agent from the inside is not merely one whose correctness condition implies that he himself is the agent of the Φ -ing in question; further, the experience *represents this as being so*, as part of its representational content (Peacocke 2003, 103-4, italics added).

experience. In particular, it cannot accommodate the sense in which the agent seems *absent* from the experience. Mylopoulos (2017) is unique among recent representationalists about agentive phenomenology in having addressed this worry. She proposes that in immersion, a first-person agentive representation still occurs (as it must on her view in order for the agent to experience herself acting), but it occurs without the agent being conscious of it (Mylopoulos 2017, 555). That is, when subjects are immersed in some activity, agentive representations are formed in the normal way so as to realize an experience of acting. But the subject does not notice its content because she is not introspecting the experience. For example, if I am immersed in performing a drop shot, on Mylopoulos's view, I may have the thought that I am performing a drop shot, but I am unaware of this thought. So, on Mylopoulos' view, my experience of performing a drop shot is an immersive rather than reflectively detached one.

However, we should resist Mylopoulos's proposed way of accommodating immersion experiences. Most significantly, her proposal seems to conflate the subjective recessiveness of *the agent* in immersive action with the subjective recessiveness of *acting*. Recall the original intuition behind the immersed/detached distinction. This was that although the agent and her internal states might figure prominently in a reflectively detached experience of acting, the agent and her internal states seem, in some sense to be explained, absent from more immersive experiences. By contrast, Mylopoulos's proposal requires that the entire experience of acting becomes subjectively vivid for the subject only when she steps back and introspects that experience. That is a less plausible claim than the earlier one. For one, the retrospective reports of flow states are of experiences of acting that *were subjectively vivid* for the subject when they originally occurred. What was (somehow) subjectively remote was the experience of *oneself as agent*. Relatedly, Mylopoulos' proposal in terms of introspection would seem to predict that the agent's sense of engagement in the activity should tend to *increase* or *intensify* as she comes to reflect more and more on her performance compared to when she was unreflectively absorbed in the act. After all, it is with introspection, on Mylopoulos's view, that the content of the experience of acting becomes conscious. But this prediction seems false on its face: famously, explicit reflection on one's own performance tends to interfere with one's sense of engagement in a task, potentially leading one to "choke".¹⁶

¹⁶This appears to be the grain of truth behind "self-focus" or "explicit monitoring" theories of choking under pressure (e.g., Beilock et al. 2004). According to these accounts, choking under pressure tends to result from explicit attentive monitoring of one's performance (e.g., the execution details of highly practiced actions). However, as Montero (2016) discusses, the correct observation that *certain forms* of conscious attention during an expert performance impair one's performance (and one's sense of subjective engagement in it) has sometimes been mistaken, both in empirical work and elsewhere, as evidence for the stronger claim that *all* forms of conscious attention to an expert performance are detrimental. I agree with Montero that the evidence cited in support of explicit monitoring theories falls far short of supporting the latter claim (and

For the above reasons, Mylopoulos's attempt to explain basic subjective action ownership in terms of occurrences of the first-person indexical is unlikely to respect the distinction between immersive and more reflectively detached agentive experiences. Mylopoulos's is just one account, but we can expect similar difficulties to arise for *any* purely representationalist account of agentive experience. To see why, it is useful to distinguish two ways to understand what it is to act or to control one's behaviour *consciously*. On one understanding of these phrases, for an agent to act consciously is for the agent to perform an action *of which* she is conscious. Here, the agent's action is the intentional object of a monitoring state that functions to register or to track such things as whether the agent is acting, what the agent is doing, and the agent's manner of doing it. Similarly, to exercise conscious control, on this understanding, is to engage in a process of control *of which* one is conscious. This is the meaning of "conscious control" and "conscious action" that representationalist accounts like Mylopoulos's target. However, this is different from the sense these expressions seem to have when used to describe immersion experiences. Here, we are confronted with an action or a process of control *to which* the agent's own conscious perspective seems actively to contribute. To explain the latter sense of "conscious control" or "conscious action", we would need to understand the basis of the subject's experience of her own subjective perspective as functionally integrated within the process of control rather than as witnessing that process unfold from without. Insofar as they provide no account of this more basic experience of oneself as practically engaged in an activity, representationalist accounts of the experience of acting falsify the phenomenology of immersive action. And in particular, they falsify the phenomenological status of the agent within these experiences of acting by suggesting that the agent is always explicitly represented in the content of the experience. In short, representationalist accounts yield an excessively detached conception of agentive experience.¹⁷

Let's turn, finally, to my solution to the puzzle about how the agent figures in the immersive experience of acting. My solution rests on the suggestion that the phenomenal character of an immersive experience of acting possesses a *perspectival character*. More specifically, on this proposal the sense of first-person ownership for an immersive action consists in a perspectival feature of the experience of acting related to the subjective or-

also that there is compelling evidence to the contrary—e.g., Eysenck et al. (2007); Eysenck and Derakshan (2011); Vine et al. (2013)).

¹⁷I include Shepherd (2017a) as one of the targets of this critique. While Shepherd has made a valuable contribution by emphasizing the constitutive role of the subject's engaged perceptual perspective in the experience of acting, his background commitment to representationalism prevents him from adequately accommodating this insight. The experience of acting, on Shepherd's view, as on Mylopoulos's, remains one of *passively informing* the agent about what she is actively doing. If so, then no explanatory significance is given to the fact that the subject's conscious perspective in an immersion experience is itself an active or controlling state.

ganization of the agent's *conscious attention*. We can think of this on analogy with how the egocentric character of a visuospatial experience is able to convey a sense of the perceiver's spatial presence within a scene without having to explicitly represent the perceiver as such.¹⁸ In a similar way, I suggest, the agent does not manifest in an immersion experience as any entity within the experience that conscious attention targets. Rather, the agent manifests as the implicit *origin* of conscious attention during the immersive task performance. By understanding the agent's subjective status as a perspectival feature of agentive experience, we simultaneously accommodate the agent's felt *absence* from immersion experience—namely, as an explicit object of conscious attention—and also the agent's felt *presence*—namely, as the origin of conscious attention. And in so doing, we solve our puzzle about the agent's subjective status in immersion experience. As it identifies the agent's subjective presence in immersive experiences of acting with the perspectival origin of conscious attention, I'll call this an "attentional account" of basic subjective ownership of action.¹⁹

The notion of "perspective" as it applies to immersive agentive experiences calls for elucidation. The type of conscious perspective agents bear in virtue of enjoying a first-person immersive agentive experience plausibly isn't *identical* to the kind of perspective subjects possess in virtue of enjoying a visuospatial perceptual experience. Notably, whereas "origin" has a literal geometric meaning when used to characterize a perceiver's felt location within a visuospatial experience, the same term has a non-geometric meaning when used to describe how the agent figures within experiences of immersed activity. The way the agent figures subjectively in an experience of acting is no more than analogous to how the perceiver manifests as located at the egocentric origin of a visuospatial experience. At the same time, however, it would be wrong to conclude from these genuine differences that an agent's first-person perspective in immersive experiences of acting is wholly unrelated to the agent's first-person *perceptual* perspective on the world. To draw that conclusion would be to miss something important about what makes the

¹⁸Cf. Campbell (2002), who holds that we do not need to use 'relational' egocentric notions (e.g., 'x is above me' or 'y is to the right of me') to state the content of a visual spatial experience, but only simpler 'monadic' egocentric notions (e.g., 'x is above' and 'y is to the right'). Ordinary human vision, Campbell suggests, "represents things as 'to the right' or 'above' using the monadic egocentric notions, rather than the relational terms" (2002, p. 184; see also Campbell 1994, 129-129). Similarly, Perry denies he is explicitly represented anywhere in his visual experience (at least in virtue of his experience's spatial perspectival character): "I am not in the field of vision: no component of my visual experience is a perception of me. How then can this experience provide me with information about how objects are related to me?" (1993, 205). Peacocke (2003) and Schwenkler (2014) offer dissenting voices in this context.

¹⁹Marcel (2003, 84 ff.) also proposes to understand the immersive sense of action ownership as a "perspectival" feature of experiences of acting. But whereas Marcel unpacks the notion of perspective involved in immersion in terms of the egocentric spatial frame of reference implicit in the motor specifications governing movement execution, I do so in terms of the motivational framing of conscious attention. See Peacocke (2003, 97-8) for a convincing rebuttal of Marcel's proposal in terms of egocentric spatial frames of reference.

phenomenology *immersive*. We must say more to unpack these claims.

I propose the following way forward. I will use what we know already about attention from Chapters 1–2 to cast light on what an agent’s “first-person perspective” in an immersion experience means. If I am successful here, I will have done two things. First, I will have explained the sense of basic ownership of action in terms of the perspectival structure of an immersive experience of acting. Second, I will have explained the latter property of agentive experiences in terms of the actual structure of attention as an exercise of conscious control. I’ll now turn to the second of these aims.

3.4 The perspectival structure of an experience of acting is based in the causal structure of conscious attention

As discussed in Chapter 1, any episode of attention can be understood as a psychological process constituted by two types of mental state: an attentional source state and an attentional resultant state (see §1.3). In this section, I want to suggest that we can use the source-resultant structure of attention in order to make explicit the sense in which the agent manifests as the “origin” of conscious attention in an immersive experience of acting. I’ll start by reviewing the source-resultant structure of attention.

An “attentional source state” is the mental state that sets the goal for an episode of attention and also directs its course. In Chapter 2, I argued that we should understand the attentional source state for an episode of automatic attention as played by the agent’s total motivational state and proximately mediated by the agent’s integrated priority map (§2.5.1). The “resultant state” of attention refers to the biasing *effect* or *result* of the source state on the mental states upon which the agent is relying to execute the goal that the source state sets. These are nonconceptual perceptual and perceptuomotor states in the case of perceptual attention. For example, in visual attention, the resultant state will be a visual state that has been modulated by the goal-setting source state. This is a visual state in which the subject is preferentially sensitive to visual information registered as of high priority according to the priority rankings of the integrated priority map. The subject is also disposed to respond to this information, when detected, in ways conducive to attaining the goal that the source state sets. Whereas the psychological process of visual attention is the biasing of visual states by source states, a psychological *state* of visual attention is the resultant biased visual state (whose status as attention partly consists in its having been biased by the source state). The causal structure of attention is therefore comprised of both motivational states (i.e. states of the agent’s priority system) and lower level implementational states (e.g., sensorimotor states).

I argued in Chapter 1 that the process of attending realizes a central part of the *guidance control* that agents manifest in skilled bodily action. Through the twin process of enhancement and suppression, agents cope with the obstacle to successfully coordinated action that distractor interference poses. I am assuming that the actions performed in subjective immersion are no exception to this claim. Indeed, subjectively immersive actions arguably provide paradigm examples of attentional guidance. Regarding the attentional source state for immersive actions, I will assume that the agent's attention is governed by the agent's total priority system. In this sense, I take it that we are dealing with a kind of automatic attention in action. However, the objectively heightened control exhibited, together with the subjectively immersive character of the accompanying experience, suggests that a single behavioural goal tends to predominate and stably govern attention and action. I assume that the guiding goal in these cases is the agent's intention to complete a certain task. Consequently, when discussing immersive actions in particular, it will not be misleading to speak simply of the agent's intention (or task set) as source state. We will return to more complex examples later.

When we actively guide our conduct, I assume that we do so, in part at least, on the basis of phenomenally conscious states and processes—i.e., conscious states and processes that “there is something that it's like for the subject” to be in (Nagel 1974). And when the agent exercises guidance partly on the basis of conscious states and processes, I will say that the agent deploys “conscious control” over her behaviour. This is consistent with allowing that there are nonconscious contributions to exercises of motor control. It is also consistent with there being forms of unconscious attention recruited in such control (e.g., Milner and Goodale 1995/2006, 204-5; Kentridge 2011; Wu 2011b, 71; Mole 2014). It is simply to affirm that conscious states and processes sometimes directly contribute to aspects of motor control, if only to relatively coarse-grained action parameters (see Briscoe 2009; Briscoe and Schwenkler 2015; Shepherd 2016a,b). One form of conscious control is conscious attentional control. By “conscious attentional control”, I mean that states and processes of conscious attention are among the conscious states and processes that agents deploy in guiding behaviour. Finally, I understand the notion of “conscious attention” as an attentional resultant state that is phenomenally conscious and whose phenomenology is partly the product of attending. For example, a state of conscious visual attention is, on my understanding, simply a conscious visual state whose phenomenology has been shaped through a process of motivational biasing. The characteristic phenomenal contribution of attending to conscious visual states is frequently expressed using various spatial metaphors, including that of “foreground” vs. “background” (e.g., Jennings 2015) and “center” vs. “periphery” (Watzl 2017). In each case, the former term is meant to connote what is experientially in the focus of one's attention, and the latter what is suppressed.

In the next chapter, I examine this characteristic subjective organization of conscious attention in more detail and examine alternative ways of conceiving of this subjective organization. Till then, I appeal to it merely as a placeholder for the phenomenological contribution that the process of conscious attention might make to an action-guiding visual experience and that qualifies it as a state of conscious visual attention.

With these materials in place, let's return to the first-personal structure of an immersive experience of acting. If, as I've suggested, conscious attention is sufficient for a type of subjective perspective, then we can pose two questions about this perspective: what is it a perspective *from*—i.e., what is its “origin”—and what is it a perspective *on* or *of*—i.e., what occupies the “field of view”? We can now use the source-resultant profile of conscious attention to answer these questions.

To begin with, the “origin” of the subject's motivational perspective corresponds to the attentional source state—e.g., the agent's current priority of completing a certain task. In saying this, we accommodate well-known properties of the immersion experience. For example, we respect the familiar observation that “flow tends to occur when the activity one engages in contains a clear set of goals” and that the specific *importance* of having clear goals (at least for engendering a sense of flow) “lies in their capacity to structure experience by channeling attention” (Csikszentmihalyi et al. 2005, 601). In an experience of engaged activity, one's conscious attention emanates from a particular set of goals, which serves to frame and structure one's attention in the situation at hand. Metaphorically, agents consciously attend in immersion from the perspective of an attention-structuring motive or set of goals, including those specified by the agent's intention.

If agents consciously attend from the perspective of a clear set of goals, then what is that perspective a perspective on or of? I take this to be a question about the *objects* of conscious attention. In other words, it is a question about what, in attending, agents attend *to*. Here, some caution is required. It's often convenient to speak loosely about the objects of attention in terms of the goal that the agent is deploying attention to accomplish, and I've sometimes allowed myself to speak this way. For example, if Sanjeev is immersed in a squash match, it may be natural to describe Sanjeev as attending to the task of winning the match. More generally, it can be convenient to speak of the object of an agent's attention in terms of the fulfillment condition that is specified by the attentional source state—e.g., by Sanjeev's current task set. This is a natural way to interpret task-based accounts of attention. According to these accounts, the “theoretically fundamental” concept of an attentional object is the task that the agent is working to achieve (see e.g., Mole 2011, 72-3; Koralus 2014, 37). If we interpret “task” to mean the goal state that is specified by the agent's task set, then task-based accounts yield the verdict that agents attend to their goals in acting—e.g., to the goal of winning a squash match or, to use an example

of Koralus's, the goal of finding an incompetently played part in an orchestral performance. Put in terms of the notion of motivational perspective, this generates the result that conscious attention realizes a perspective on a task or goal.

But that seems incorrect. The agent's goals (including those specified by the agent's task set) feed into the "subject side" of an episode of conscious attention. They are our point of view or frame of reference. They are not what we take that perspective on or toward. When Sanjeev is immersed in the task of winning the match, the task that guides his conscious attention could itself be secluded from his conscious attention. Unless one is strategizing about what to do next, it may even be that attending to one's goals will tend to distract one from successfully acting on them. For example, devoting too much attention to certain incentives may be detrimental to spontaneity. Consequently, we should not regard the agent's goal or task as providing the object or objects of her attention. Rather, I propose to understand the object or objects of conscious attention in terms of the entities figuring in the contents of attentional resultant states (e.g., perceptual states). On a resultant-based conception, the objects of attention include the objects of perception (in the case of perceptual attention) and of thought (in the case of intellectual attention). For example, when Sanjeev is attentionally engrossed in a squash match, his focus will be on such things as the changing position of his opponent and the ball, his own grip and foot placement, and his current distance from the "T".²⁰ Unless he becomes distracted, he won't be attending to the conversation occurring outside the court or to the colour of his opponent's watch or socks. At least, he will be attending to them *less* than he is to other things. On this proposal, a statement like "Sanjeev is attending to the task of winning the match" is therefore shorthand for a statement describing the selective organization that Sanjeev's conscious perceptual states have acquired as a result of top-down biasing by his intention to win. "Attention to a task" amounts to, for example, patterns of perceptual attention to the world structured according to the agent's goals.²¹

²⁰I limit myself to the objects of Sanjeev's conscious perceptual attention. A fuller discussion would include the objects of his intellectual attention (e.g., the contents of strategic planning and decision-making). I assume that there is a parallel sense in which Sanjeev's intellectual attention is directed toward the world rather than explicitly toward his goals.

²¹My view here is partly inspired by the following passage from Alan White (1964):

Whether or not a man is giving attention to his fishing depends on whether he looks at the water, listens to the rustling of the reeds and thinks about the habits of fish or whether he looks at his watch, listens to the aeroplane in the sky and worries about his children's education and his stocks and shares. The attentive is distinguished from the inattentive fisherman by the relevance or irrelevance to his fishing of what he gives his attention to ... To be attending as an agent to what one is doing ... signifies attention, by way of looking, listening, thinking or doing to various objects relevant to the general task attended to and absence of attention, in such ways, to objects not relevant to the general task. (16)

Here, White gestures at a way to translate statements about a fisherman's attention to a task (e.g., fishing) into statements about attention to various "objects" (e.g., objects of perception and thought) that are rele-

I've suggested that an agent's motivational perspective within a first-person experience of acting is a matter of how the agent is consciously attending in acting. I've further suggested that the objects of the agent's conscious attention partly coincide with the objects of perception. It follows that an agent's motivational perspective in a first-person experience of acting partly coincides with her perspective on the objects of perception, including her own intentional movements. These two perspectives only partly coincide for two reasons. First, motivational states structure cognitive states in addition to perceptual states. So, motivational perspective encompasses more than perceptual perspective. Second, there are elements of a perceptual perspective that are not part of motivational perspective (e.g., egocentric spatial perspective). Still, motivational and perceptual perspectives genuinely intersect where conscious perceptual states are motivationally structured through the process of conscious attention. Adopting a motivational perspective on one's action constitutively endows that action with the basic feeling of being one's own.

Whichever way we choose to characterize the objects of the agent's conscious attention in immersion, what is crucial is the way conscious attention is being directed. Whether we describe agents as attending to the task at hand or instead to task-relevant information, the experience of basic ownership of an action depends on the practical use to which conscious attention is put in action—i.e., in service of goal execution. Specifically, my view is that the agent will experience basic ownership over an action if the motivational state guiding her behaviour is also the source state of her conscious attention. Put more metaphorically, the agent feels basic ownership toward an action if her motivational perspective is engaged in guiding action, where "motivational perspective" refers to the motivational state that occupies the functional role of attentional source state. I'll conclude this section with a few clarifications about the proposal.

First, I mean this to be a supervenience thesis. In particular, there can be no change in the experience of action ownership without a corresponding change in how the agent exercises her conscious attention in action. I am suggesting that facts about the way an agent is exercising her conscious attention in action necessitate facts about the phenomenology of subjective ownership. My main concern here is to defend this as a thesis about the actual world. In particular, in §3.6 I will respond to potential empirical challenges that one might raise against the claim that the phenomenal property of action ownership does in fact depend on exercises of attentional control. Against this, I will suggest that my necessity claim is consistent with the known empirical conditions under which subjective

vant to this task or goal. In particular, White seems to regard attending to a task as a less basic notion than attending (as a "spectator") to the world around oneself, where spectator attention is harnessed motivationally by one's task. For complementary points, see Evans (1970, 109 ff.) and Roessler (2000; 2003, 389-90) for helpful commentary on Evans (1970). For an opposing task-based development on White's views, see Mole (2011).

ownership is lost.

Second, while I'm only defending a supervenience claim for action ownership experience, the condition explains salient features of a wide range of cases of agentive experience. First and foremost, it captures that, in ordinary cases of skilled bodily action, the agent's conscious attention is *motivationally structured*. The motivational state that guides the agent's action is also the source state governing how conscious attention is allocated. And it is this fact—that the behaviour-guiding motivational state occupies the role of conscious attentional source—that I take to explain why the agent experiences herself as engaged in the task at hand.

Third, while it highlights that, in ordinary bodily action, conscious attention is harnessed in service of behavioural guidance, my account does not overstate conscious attention's behaviour-guiding role. In particular, it does not imply that felt ownership extends *only* to actions or properties of actions that conscious attention guides. This is important because there are contexts in which an agent feels ownership over an action or properties of an action, but the agent plausibly doesn't use her conscious attention to guide that action or those properties of the action. Two such cases warrant mention.²²

First, it could turn out that agents experience ownership toward properties of their actions that, as a matter of empirical fact, are settled outside conscious attention. Fine-grained properties like precise finger placement during a rapid, highly practiced routine might be like this (though, as discussed in Chapter 1 (§1.4.2.1), it is not a straightforward matter determining whether the states responsible for fine-grained execution details are conscious or not). While on my view these properties of action are the product of attention (see §1.4.2.3), they might not be the product of *conscious* attention. My proposal leaves open the possibility that there are properties of the agent's behaviour to which conscious attention itself contributes little or nothing, but that are still subjectively owned. This will be the case if the motivational state that is responsible for settling those fine behavioural details is the same motivational state that occupies the source of one's conscious attention. In that case, it is still one's motivational perspective that guides one's fine-grained movements. Consequently, my proposal leaves open exactly to what degree a subjectively owned motivational state guides behaviour in virtue of harnessing conscious attention. What matters for establishing sense of ownership over some property of one's action is not whether conscious attention settles that property but whether it is settled under the guidance of the source state of one's conscious attention. My proposal in this chapter is thus consistent with the possibility that agents experience ownership over even unconsciously programmed properties of their behaviour.

To take another example in which conscious attention itself arguably does not play an

²²I am indebted to Myrto Mylopoulos for raising versions of the following points in discussion.

action-guiding role, consider the mental act of consciously attending to something. Plausibly, conscious attention is (or can be) an action over which one feels ownership, as when one experiences oneself actively holding one's attention on a body part. But, one might think, it's highly implausible to suppose that acts of conscious attention *themselves* require guidance by a separate act of conscious attention in order to be subjectively owned. My proposal does not suggest otherwise. Subjective ownership of an action, on my proposal, does not require that conscious attention itself guide the action. Rather, it requires that the motivational state guiding the action also be the source of one's conscious attention. This requirement trivially obtains in cases of active attending, since necessarily, any motivational state that directs an act of conscious attending thereby serves as the source of conscious attention. Consequently, my proposal not only allows, but actually requires that all acts of conscious attention be subjectively owned. Put differently, if conscious attention is an action, then on the present account it is, apparently unlike other actions, essentially subjectively owned. If conscious attention is an action, it is the agent's act of taking up a motivational perspective.

The consequence that acts of attention are uniquely subjectively inalienable might give us pause. In particular, one might think that the case for the attentional account will not be secure until it receives independent defense. I take up this challenge in Chapter 4. There, I will show that an independently motivated view of how conscious attention perspectivally structures an agent's perceptual awareness of the world requires that we understand episodes of conscious attention as being necessarily subjectively owned. Till then, I rest my case for this claim on its intuitive plausibility.

This concludes my account of basic subjective ownership of immersive actions. We can summarize the proposal in terms of two steps. In the first step, I explained the phenomenal property of basic subjective ownership in immersive experiences of acting in terms of the "perspectival character" of immersion experiences. According to that proposal, the agent is present subjectively in an immersive experience of acting in something like the way a perceiver is subjectively located at the egocentric origin of a visuospatial experience. In the second step, I explained what "first-person perspective" amounts to in an immersive experience of acting by reference to the causal structure of attention as a form of conscious control. According to this proposal, the agent manifests as the "origin" of an experience of acting in the sense of occupying the source state of conscious attention in action execution. I turn next to show how we might extend this account of the first-personal element of immersion experiences to various non-immersive experiences of acting.

3.5 Extension to non-immersive experiences of acting

Not all experiences of acting are immersive. Nevertheless, motivational perspective may characterize all first-person experiences of acting, even if less vividly. In this section, I extend the account of subjective ownership that I've given for immersion to paradigmatically non-immersive agentive experiences, including experiences of acting half-heartedly, of acting anxiously, and of acting with self-control.

One way to alter the immersive character of an agentive experience is to vary the strength of the agent's motivation to perform the action. To begin with the clearest sort of case, actions performed with only weak motivation are unlikely to be subjectively immersive. To take an example of weakly motivated action from Shepherd (2017b), consider the act of grading papers when you want to be doing something else, like playing with your dog. Your action of grading is intentional, but, let us suppose, weakly motivated or halfhearted—e.g., you are not marshaling special effort toward the task. Shepherd argues that acting halfheartedly (i.e., with weak overall motivation) decreases the control one possesses over an action. For example, supposing one has already initiated Φ -ing, then weak motivation to Φ leads to decreased guidance over Φ -ing. And as he notes, performance costs associated with weak motivational levels are widely taken to be mediated by increased distractibility (e.g., "goal-neglect"). Weak motivation to Φ decreases the agent's guidance over Φ -ing by decreasing focus on information pertinent to Φ -ing (2017b, 268-9).

To these points about attentional guidance control, I add a claim about how the phenomenology of action-ownership. Specifically, I propose that the same variable underlying decreased guidance control in halfhearted action—i.e., increased distractibility—also underlies its non-immersive phenomenology. When motivation to Φ is weak, motivational factors other than one's intention to Φ come partly to occupy the source of conscious attention, sometimes diverting attention from Φ -relevant information toward Φ -irrelevant information. Since attention is central to action-guidance, goal-attainment will be more due to luck under conditions of weak motivation than when attention controls performance. Action success will also be less *up to you* (or less *your doing*) and will phenomenally manifest as such. You are *less engaged*.

For analogous reasons, it's possible that being too strongly motivated to act also disrupts immersion. Specifically, if beyond a certain threshold, motivational strength induces anxiety about failure, then the relaxed focus of immersion will give way to a more anxious performance. Conscious attention is apt to become diverted from intention-relevant information to intention-irrelevant information. For example, attention may turn (counter-productively) to well-rehearsed execution details, toward hypothetical negative outcomes, or to possibly hostile faces in the audience. If so, then in conditions of both

weak and (very) strong motivation, one is likely to lose experience of being immersed in a single action. Under both conditions, one's motivational perspective partly loses its internal coherence (because jointly occupied by interfering motives), and behaviour will be less coherent and smoothly guided than in immersion.

So far, I've been discussing forms of ownership experiences that could be enjoyed even by an unreflective "wanton": an agent whose behaviour is exclusively guided by first-order desires and who never considers the question of how they want to be motivated. However, as discussed in §3.2.2, philosophers of action have sometimes been concerned with loftier forms of agency ("agency par excellence" or "autonomous agency"). As we saw, one theoretical challenge, in this context, is to understand how agents achieve autonomy in the face of competing motivational states, biasing the outcome in favour of a self-consciously endorsed motive. Does the attentional control account of subjective ownership have anything to say about the self-reflective ownership associated with self-controlled agency?

It may. Autonomous ownership may be a special case of conscious attentional control. In this context, it's illuminating to consider a constraint Velleman (1992) posits regarding the psychological state whose role in behavioural is constitutive of the agent's participation. Against Frankfurt's suggestion that the agent can be reduced to higher-order motivational states, Velleman objects that merely being of a higher-order isn't sufficient to occupy the role of the agent, since a motive of a higher-order can also come under reflective scrutiny and, in principle, be reflectively disavowed. Regarding the agent-constituting mental state, Velleman writes:

[T]he reason why it cannot be played by anything that might undergo the process of critical review is precisely that it must be played by *whatever directs that process*. The agent, in his capacity as agent, is that party who is always behind, and never in front of, the lens of critical reflection, no matter where in the hierarchy of motives it turns. (Velleman 1992, 477, italics added)

Velleman's idea is that to occupy the role of the agent, a psychological state must be inalienable in the sense of always directing practical reflection without ever becoming an object of such reflection. His solution is to claim that practical reasoning is directed by the motive to act on the basis of reasons (or to make sense to oneself). Such a motive would, Velleman notes, function to direct practical reasoning, but could not itself be reflectively disavowed. This is because to call it into question would be to manifest the concern to act for reasons and so for the motive to continue to direct one's thought. Consequently, Velleman suggests, we have uncovered a motive suitable to serve as the agent: a motive that occupies one's conscious perspective when and only when one is engaged in practical

thought. Because it is a motivational state, moreover, its contribution to practical reasoning may be sufficient to tip the balance in favour of behavioural tendencies that, without it, would not be strong enough to prevail (e.g., the desire to not use a drug). Hence, it would also be able to solve Velleman's puzzle around self-control (see §3.2.2).

While we may dispute the details of Velleman's proposal, the constraint that he proposes in the above passage seems to me fundamentally right. Moreover, it seems to imply commitment to a version of the attentional account of subjective ownership. In particular, Velleman assumes that the agent figures in practical thought not as a potential object of conscious (in this case, intellectual) attention, but essentially as that which directs conscious attention. However, Velleman seems to assume that we must individuate the agent-constituting state by its distinctive content or satisfaction condition. And it is this assumption that leads Velleman to search for a motivational state that, given its content, cannot but figure as the agent's mental perspective when thinking practically. I reject that assumption. My alternative is to say that the agent-constituting mental state is not individuated by a distinctive content or satisfaction condition. It is simply individuated by the functional role of directing conscious attention. This allows us to retain the assumption that the agent-constituting state is what directs processes of conscious attention. But we can allow that states with different satisfaction conditions will tend to play this attention-directing role on different occasions: each with an equal claim to fulfilling the role of the agent (when it is playing this role). In special contexts of self-control, it may be that a motivational state with particular satisfaction condition comes to direct conscious attention—perhaps the motive to act for superior reasons. This is an important and distinctive agential capacity, and one that can be made sense of on the present account. However, it is also just one of many motives that can direct conscious attention. We do not need to single it out as the unique setting in which the agent is genuinely active.

Having indicated how the attentional account might be extended to explain the sense of ownership for non-immersive agential experiences, there is hope for a fully general account of the first-personal aspect of agential experience in terms of motivated deployments of conscious attention.

3.6 Losing ownership?

In this final section, I will explain how my proposal accommodates intentional behaviours that are abnormal with respect to subjective ownership and also identify some consequences of my proposal for the relation between consciousness and agency. I will use ("thetic") representationalism about agential experience as my main foil, as it most clearly embodies assumptions that I take myself to be rejecting. It should be noted, how-

ever, that my aim in this section is not to argue against representationalist accounts (which I did in §3.3), but to further elucidate where our views differ and also to fend off a potential line of objection to my position arising from nonstandard experiences of acting.

Recall that the attentional account in this chapter commits to a supervenience claim: there can be no change in the experience of action ownership without a corresponding change in the control the agent exercises—specifically, in how the agent deploys conscious attention to the action being performed. Consequently, the account will be in trouble if it is possible for a pair of actions to differ in subjective ownership but not differ in the conscious attentional control the agent exercises in executing the action.

By contrast, representationalism is consistent with the possibility that subjective ownership fails to supervene on action control. On that view, an experience of acting is realized by a mental representation whose role is to monitor the agent's own actions, and the phenomenal character of the experience of acting is explained by the representational content of the monitoring representation. This view can allow for pairs of cases in which action control remains fixed, but in which subjective action ownership differs. It can do so because it can allow that how the action is represented in the two cases differs. For example, there is no barrier, on representationalism, to an action that is normally subjectively owned occurring without agentic experience at all. This would be expected to occur, on their view, if control were exercised without an accompanying agentic representation. Just as the same type of object or event that is perceptually experienced at one moment can also exist or occur unperceived at another moment, similarly the representationalist about agentic experience can accommodate the possibility that the same type of agential control that is at one time accompanied by a first-person experience of acting can at another time occur wholly unconsciously.

To take another example, the representationalist can allow for the possibility that an agent performs an action and even experiences what she is doing *as* an action, but does not experience *herself as the agent* of the experienced action. For example, the representationalist can accommodate the possibility that the monitoring state that realizes an experience of action misrepresents the owner of the experienced action—e.g., misattributing the represented action to an external source. Thus, they have their own way of accommodating the extreme examples of “alienation” or “passivity experience” discussed in §3.2.3. Just as on a standard representationalist view of perceptual experience, a perceptual experience can be illusory in virtue of misrepresenting the properties of a perceived thing or event, similarly a representationalist about agentic experience can accommodate the possibility of an agentic experience that inaccurately represents who is acting. On such a view, it is possible that the only factor distinguishing a first-person experience of acting, an unconscious action, and an alienation experience is how the action is represented as

being. There need be no difference at the underlying level of action control (e.g., in the causal basis of the agent's performance).

It might seem that the apparent possibility of alienation experiences and unconscious action is evidence *for* representationalism and *against* my attentional account of subjective action ownership. In particular, such cases might appear to provide evidence against the claim that action control processes necessitate agentic experience. As I'll now show, however, the attentional account has its own way of accommodating these putative problem cases. I'll suggest that, in each case, there is a difference at the level of action control—specifically, in deployments of conscious attention—that can explain the observed differences in agentic experience. If this is right, then there are not clear empirical grounds favouring one of the views over the other, at least with respect to the class of phenomena under consideration. The case for deciding between the two views instead lies elsewhere.

Let's begin with alienation experiences: experiences in which a kind of motivational control over behaviour is exercised without any felt ownership. In particular, let's begin with the hypothetical examples of subjectively passive motivational state-causation Horgan and colleagues use in order to motivate the existence of "self-as-source" phenomenology within ordinary experiences of acting (see §3.2.3). I suggest that we can explain the subjectively passive character of these experiences by grasping how the agent's conscious attention is allocated while her motivational state is being executed. Recall that on the attentional account, when the conscious motivational state directing an agent's behaviour also directs conscious attention (i.e., occupies the attentional source role), the agent experiences the behaviour as her own doing. Now, consider that the alienated subject in Horgan's cases is described as "witnessing", "watching", or "observing" her own motivational state as it generates the motivated behaviour (e.g., observing her urge to clench her fist passively causing her fist to clench). Arguably, these acts of reflective self-observation—acts of conscious attention—are not *themselves* motivated by the state that the subject is observing—namely, the urge. That is, in reflecting on a motive (in the way Horgan invites us to imagine), the subject's self-reflective attention is not contributing to the satisfaction of the motive that is being reflected upon. If so, then the behaviour-guiding motivational state (e.g., the urge to clench) fails to occupy the *source* of the subject's conscious attention. Consequently, on the attentional account, it doesn't manifest as the agent's motivational perspective. So, the type of conscious control that, on the attentional account, explains the agent's felt presence in the experience of acting is manifestly absent in the sorts of cases Horgan and colleagues describe.

The conceivability of passive motivational state-causation does not show that when agents have normal experience of acting, they experience their behaviour as controlled by

something other than a motivational state.²³ And it certainly does not present evidence that in order to accommodate self-as-source phenomenology we must include an explicit representation of the agent, as such, within the content of a “thetic” or “mind-to-world” mental representation—e.g., <I am clenching my fist>. As I argued in §3.3, that would *itself* generate an excessively detached conception of agentive experience.²⁴ Rather, the conceivability of passive motivational state-causation suggests that the presence or absence of basic action ownership depends on how the agent’s conscious attention is being deployed during the implementation of motivated behaviour. In particular, it depends on whether or not conscious attention is harnessed toward that motivational state’s fulfillment.

The attentional account thus has its own explanation of the alienated phenomenology Horgan’s cases vividly illustrate. Let us turn, then, to consider real-life alienation experiences, like those present in extreme cases of schizophrenia. Since these patients suffer alienation experiences while intentionally moving their bodies, one might think these cases establish that subjective action ownership fails to supervene on exercises of conscious control. Indeed, I agree that alien control experiences are evidence for a *kind* of conscious control without subjective action ownership. As I emphasized when considering Searle’s account (§3.2.3), such cases suggest that a conscious intention can control behaviour without the subject experiencing action ownership. However, to refute the supervenience claim, one must claim not only that there is conscious control exemplified during an alienated experience of acting, but that the conscious control exercised is the same as that exercised in the first-person experience of acting. And the empirical support for the latter claim is weak.

To begin with, it is now well-established that schizophrenic patients exhibit significant attentional disfunctions. In a review of current research on a neural system sometimes

²³Here, we would do well to distinguish the claim that (i) in a normal experience of acting, one does *not* experience one’s behaviour as controlled by a motivational state, and (ii) in a normal experience of acting, one experiences one’s behaviour as *not* controlled by a motivational state. Claim (ii) strikes me as less well-supported than (i) both introspectively and from reflection on the cases Horgan and colleagues advance. Notably, though, the attentional account is not only consistent with claim (i) but predicts it. If in a normal experience of acting, the behaviour-controlling motivational state directs conscious attention (i.e., occupies the role of attentional source), then this state will tend to remain secluded from conscious attention, leaving determinate information about it inaccessible introspectively.

²⁴Admittedly, it’s not fully clear if this is the moral Horgan himself believes we should draw from his reflections on passive phenomenology of motivational-state causation. In most contexts, Horgan writes as though this *is* the moral we should draw. He seems to suggest that we can account for self-as-source phenomenology by explicitly including a representation of the agent, as such, within the experience’s veridicality conditions (e.g., <I am Φ -ing >). Elsewhere, Horgan says things more congenial to the type of account I’m proposing—in particular, by denying that we should explain self-as-source phenomenology with an explicit “I” representation (see e.g., Horgan and Nichols 2015). Whatever Horgan’s own views on the matter, others (e.g., Bayne & Levy 2006 and Bayne 2008) have clearly wished to draw representationalist morals from Horgan et al.’s observations.

called the “motivated attention system”—an integrated collection of neural regions underlying the general “ability to direct attention to stimuli relevant to the current goal and respond to those stimuli while ignoring competing responses”—Liddle (2006) concludes that “recruitment of this motivated attention system is impaired in schizophrenia.” (196). Specifically, “patients with schizophrenia devote greater resources to processing novel stimuli, even though these stimuli are irrelevant to the current task, implying a decreased ability to allocate resources in accord with current goals” (200). Liddle primarily discusses schizophrenic impairments at oddball target detection tasks, but attentional deficits in schizophrenic patients have also been discovered for other tasks, including visual search (Gold et al. 2007) and visual selective attention (Carter et al. 2010). In short, schizophrenic patients exhibit disruptions in the control of conscious attention when executing a variety of tasks: they are likely *not* exercising the same kind of cognitive control as neurotypical subjects are.

The above evidence is sufficient to show that schizophrenic alienation experiences are not counter-evidence to the supervenience of subjective ownership on deployments of conscious attention. Still, it would be desirable to see evidence for an explanatory connection between the impairments in attention and the abnormal experiences of acting. While nothing in this area is uncontroversial, there is reason to take seriously the claim that these two facts are explanatorily related.²⁵ A frequently discussed feature of schizophrenic patients exhibiting passivity experiences is decreased sensory attenuation of self-administered tactile sensations. Whereas in neurotypical patients subjective intensity (e.g., pleasantness and ticklishness) of a self-generated tactile sensation is dampened relative to externally generated ones, in schizophrenic patients suffering from passivity experiences tactile sensations are comparably subjectively intense in both self-generated and externally generated conditions (Blakemore et al. 2000). A standard explanation for schizophrenic alienation experiences is framed in these terms. Unlike neurotypical populations, who are focally aware of bodily sensations only when acted upon passively, schizophrenic patients experiencing alien control are focally aware of these sensations also while acting, making those sensory states “hypersalient”.²⁶ That is, according to this

²⁵Cf. Spence et al. (1997), who, on the basis of PET scan evidence of schizophrenic patients experiencing alien control, propose an explanation of these experiences in terms of abnormalities in the control of attention:

Given that these hyperactive cerebral regions subservise attention to internal and external bodily space, and the attribution of significance to sensory information, they provide a plausible anatomical substrate for the misattribution of internally generated acts to external entities: the cardinal feature of delusions of passivity (alien control) (1997, 1997).

²⁶Cf. Frith:

I believe we can now achieve some understanding of PH’s experiences because of what we

explanation, passivity experiences partly consist in subjectively unattenuated sensory experiences of one's own body while acting. Significantly, selective sensory attenuation (i.e., inhibition or suppression) and dis-attenuation (i.e., enhancement) are the hallmarks of conscious perceptual attention. For example, according to Treisman's classic "attenuation model" of attention, the strength of unattended stimuli is attenuated (rather than its channel completely blocked) relative to attended stimuli (Treisman 1964; cf. Mylopoulos 2012, 113). These standard claims about both schizophrenic passivity experiences and perceptual attention are consistent with an account of passivity experience as arising from a disturbance in the motivational harnessing of conscious attention. It may be that schizophrenic passivity experiences during bodily action arise, at least partly, from failure of the agent's guiding intention to harness conscious perceptual attention. This could lead to an uncharacteristically observational deployment of conscious perceptual attention during intentional movement, somewhat like the hyper-observational posture described in Horgan's hypothetical cases. The pathologically third-personal experience of acting may be based in an anomalous allocation of conscious attention during intention execution.²⁷

Whereas subjectively alienated agents consciously attend as spectators toward their own purposive bodily movement, putative cases of unconscious action ("phenomenal automatism") arguably are not performed with any conscious attention to the action. The claim, here, is not merely that aspects of the agent's action are controlled without conscious attention. That alone would not distinguish phenomenal automatism from ordinary skilled actions in which fine-grained motor parameters are programmed outside conscious attention. In the latter case, the agent plausibly experiences ownership over the entire action, even if properties of the action are specified unconsciously. I've accommodated the latter possibility with my claim that an agent feels ownership over an action so long as the same motivational state directing unconscious parameter setting also directs conscious attention. By contrast, in phenomenal automatism, conscious attention is not serving action execution *at all*. In the latter respect, phenomenal automatism is similar to alienation experience. Unlike the latter, though, during phenomenal automatism conscious attention is not deployed observationally toward one's own bodily move-

have discovered about the brain. In our normal state we are hardly aware of the sensations that occur whenever we move. This is because our brain can predict these sensations and suppress our awareness of them. But what would it be like if something went wrong with the prediction and we became aware of the sensations? Normally I am only aware of the sensations when someone else moves my hand. Such a brain abnormality could explain why PH feels as if her arm is being moved by someone else. She is abnormally aware of her bodily sensations when she moves her hand. For her it really does feel as if someone else were moving her hand. (2007, 109)

²⁷See Clark 2016, 217-19 for related suggestions within a Bayesian predictive coding context.

ments either. Rather, the agent is simply not consciously attending to what she is doing in phenomenal automatism, whether practically (as in ordinary experience of acting) or observationally (as in alienation experiences).

My aim in this section has not been to defend my account on empirical grounds. Rather, it has been to show that the supervenience claim to which the attentional account commits is sustainable in the face of alleged counterexamples. In particular, we can acknowledge the subjective differences that exist between first-person experiences of acting, alienation experiences, and unconscious actions, while maintaining that subjective action ownership supervenes on conscious attentional control. The reason that I emphasize this point is that representationalists can also acknowledge the above empirical evidence showing differences in how conscious attention is allocated in different cases of agentic experience. They will simply attribute a different significance to these findings.

For example, a representationalist is likely to interpret the hyper-observational deployments of conscious attention that schizophrenic patients engage in while acting as the effect of an antecedent alienation experience that misrepresents the agent's own action as that of another (e.g., "I am not the one doing this!"). On such a view, the schizophrenic subject represents an action as not her own doing, and this (combined with violation of expectations) motivates a shift of perceptual attention, in an observational manner, to the sensory details of how the subject's body is moving (Mylopoulos 2012, 113-4; Mylopoulos 2015, 773). By contrast, on the attentional account, there is no *antecedent* alienation experience—e.g., a mental representation as of oneself being passively moved—that causes the agent to attend observationally to how her body is moving. Rather, these observational deployments of conscious attention to movement execution *constitute* the experience of acting as being alienated rather than subjectively owned. On the latter view, what explains the active rather than passive character of the agent's experience during intentionally controlled movement is how the agent's motivational perspective is deployed in that process. Only in a first-person experience of acting is the agent's conscious attention recruited toward ensuring the success of the motivational state that is guiding bodily movement. So, both the attentional account and the representationalist have ways of understanding the available evidence. The real case against representationalism, in my view, comes from their inability to provide a satisfactory solution to the puzzle about the agent's presence in immersive experiences of acting (see §3.3).

I'd like to close by remarking on *ownership itself* on this account. This will allow me to answer an objection raised in §3.2.4 against my claim to be offering a solution to the Disappearing Agent Problem. The objection, recall, was that the Disappearing Agent Problem is a problem about the nature of action ownership, not about its phenomenology. According to the objection, we can no more grasp the relation that agents bear to their

actions through an account of how it *feels* to the agent to enter into that relation than we can explain the nature of a material thing (e.g., an apple) by offering an account of how it *looks*.

We are now in a position to grasp why this objection fails. In particular, the objection assumes that the agent's sense of ownership of some action is "of" herself as the agent of the action in the same way that, for example, a visual experience of an apple is *of* the apple one is looking at. But the attentional account reads the "of" in "sense of ownership" differently. The agent's experience of ownership does not consist in a representation of herself as an agent—e.g., "I, the subject of this experience, am the agent of this action". Neither is the experience of ownership any other contingent sign or marker of one's own agency. Rather, on the attentional account, the sense of ownership is "of" ownership in the way one's experience of seeing an apple is an experience *of seeing* (rather than of, say, hearing). That is, it is a type of conscious occurrence. If so, then just as there is something it is like for you to see the apple, so there is something it's like for you to direct conscious control toward your conduct and for that control to be your own. Most will agree that it is an error to suppose that one can grasp the nature of a conscious visual experience apart from grasping its phenomenal character. For the same reasons, it is an error to suppose that we can grasp the ownership agents bear to their actions apart from grasping *its* phenomenal character. This is because, on the attentional account, the ownership that agents bear to their conscious actions consists in the instantiation of a phenomenally conscious point of view. It consists in the instantiation of a motivational perspective. Therefore, to adequately characterize the ownership relation that holds between an agent and an action, we must grasp the constitutive role that is played by the agent's conscious perspective in making the action the agent's own.

Chapter 4

Attentional Structure

4.1 Introduction

In the previous chapter, I drew an analogy between egocentric visuospatial perspective and the subjective perspective agents possess in virtue of attending consciously (i.e., “motivational perspective”). I suggested that the agent is subjectively present in an experience of acting in something like the way a perceiver is subjectively present at the egocentric origin of a visuospatial experience. As the perceiver is subjectively present at the unrepresented origin of egocentric space, so too the agent is present subjectively at the unrepresented “source” or “origin” of conscious attention. This fact, I suggested, helps us to explain the first-personal element of an experience of acting.

Reflection on this analogy invites a corresponding question about the *structure* of a motivational perspective. Given that how the world appears to the subject is organized by a visuospatial perspective into egocentric spatial relations (e.g., of near and far, to the left and right, etc.), is there some corresponding way that acts of conscious attention subjectively organize a subject’s perceptual experience? And if there is, what is that subjective organization? In this chapter, I address these questions. In doing so, I shift my focus from how the agent figures in conscious attention to how the world appears to the agent in virtue of conscious attention.

There are two main options for an account of the subjective organization of conscious attention. The first is to regard this organization as primitive or irreducible to other phenomenal elements (e.g., Watzl 2011, 2017; Jennings 2015). Primitivists take conscious attention to endow a unique organization on perceptual experience. In their view, just as spatial perspective organizes a scene into egocentric spatial relations (e.g., of near and far), attentional perspective organizes visual experience into *sui generis* attentional relations of “central” and “peripheral” or “foreground” and “background”. On their view, the latter organization is wholly original to attention. Absent conscious attention, noth-

ing in the world or in experience would, in the relevant sense, be central or peripheral to anything else. The second option is to attempt to explain the subjective structure of conscious attention reductively in terms of aspects of perception that are already, at least partially, in place pre-attentively. Whereas primitivists regard an act of conscious attention as the absolute source of novel subjective organization, reductionists view attentional organization as an elaboration of forms of perceptual organization that do not originate with conscious attention.

This chapter defends the latter, reductive view of attentional organization. I argue that the difference conscious attention makes to the phenomenal character of perceptual experience is the product of top-down modulation by the subject's motivational states on how the world perceptually appears to the subject. On this view, attending does not contribute a totally novel subjective organization to perceptual experience. Instead, it contributes to perceptual experience by modifying and elaborating upon structures that are inherent in a visual appearance. At least with respect to the subjective structure of a perceptual appearance, conscious attention does not contribute anything original or *sui generis*.

It is important to distinguish my reductive claim about the subjective structure of conscious attention from the stronger claim that *all* of attentional phenomenology can be explained reductively. I reject the latter position. As we will see, conscious attention is a mental act. And attentional phenomenology includes a sense of ownership over one's act of attention. But on the position that I defended in the previous chapter, we cannot characterize the experience of first-person action ownership independently of conscious attention. My account of ownership phenomenology crucially invoked conscious attention as an ingredient. Since ownership of an act of conscious attention is an aspect of the phenomenology of conscious attention, and since action ownership is *itself* explained in terms of the control of conscious attention, then a fully reductive account of attentional phenomenology seems to be off the table. What remains available and, in my view, attractive is a partly reductive account of attentional phenomenology. This account is non-reductive about the sense of ownership but reductive about the subjective structure of our attentional awareness of the world. This is the account that I defend here.

Part of what is at issue in this debate is the extent of the agent's active role in perceptual experience. According to a traditional view of perception, a sharp distinction exists within consciousness between what you are passively "given" in experience and the "constructive acts" you perform on the given (Lewis 1929; Price 1932). The former refers to what is "immediately present in consciousness". The latter refers to the way that you "select from it, emphasise aspects of it, and relate it in particular and avoidable ways" to your purposes and interests (Lewis 1929, 52). Traditionally, philosophers who insist that

there is a given element in experience claim that it “is what remains unaltered, no matter what our interests” and that “no activity of thought can create or alter” it (Lewis 1929, 53). These authors intended these claims to extend to attention. For example, Price (1932) dismisses the claim (which he attributes to Bradley but reportedly finds in “many of the critics of the Given”) that “as we attend to something, this something becomes more and more ‘clear’”(1932, 16). Price finds the proposal “so extraordinary that it is hard to see how any one can have the audacity to hold it” (1932, 16). By contrast, the philosophers I will be engaging with in this chapter are ready to grant that attention alters how things appear to the subject—e.g., by making them look “more and more clear” in Price’s words. Nevertheless, even here a distinction is frequently upheld between the organization that accrues to an experience in virtue of the appearances that it passively presents to the subject and the organization that agents actively *impose* on their experiences through acts of attention. In upholding a strict distinction between the appearances that are immediately presented to the subject and the organization that the subject *adds* to these appearances through acts of attention, shades of “the given” persist into contemporary philosophical discussions of attention.

Against this, I’ll argue in this chapter that proper grasp of conscious attention’s contribution to the subjective organization of perceptual experience requires a more thorough rejection of the traditional separation between whatever is passively presented in visual experience and the active structure superimposed on appearances through attention. In rejecting primitivism and nearby views, I will be advancing a view of appearances as themselves partly the product of a motivationally guided process. On my account, attending consciously is the process whereby the agent’s motivational states guide the construction of an appearance. The subjective organization of conscious attention, as we’ll see, is the product of that process.

The plan is as follows. I begin in §4.2 by introducing some empirically established effects of attention on visual appearances, including what I call the “sharpening” and “boosting” of visual appearances. In §4.3, I consider an important recent argument from Watzl (2011, 2017, Ch. 9) and, independently, from Wu (2011c; 2014b, Ch. 4) against the claim that the phenomenology of a state of conscious visual attention consists in any property of how things *appear* to the subject of the experience. According to this argument, attentional phenomenology cannot *consist* in a type of visual appearance because, for any difference that attending might make to how things visually appear, this difference can be replicated without replicating what it’s like for the subject to consciously attend (the “replication argument”). In response, I argue that the introspective evidence on which the replication argument draws is consistent with a reductive account of the subjective structure of conscious attention. Further, I argue that the latter account is preferable to either

of the proposals Watzl and Wu advance in light of the relevant introspective evidence. In §4.4, I respond to an objection that reductive proposals like mine fail to accommodate intuitively holistic aspects of attentional phenomenology. Lastly, in §4.5, I argue that though the replication argument fails to achieve its official aim, the introspective evidence that it draws on supports a different conclusion about attentional phenomenology. In particular, it supports a conclusion about its distinctive status as a mental act. This will allow me to provide independent support for a critical, but until now undefended, consequence of the account that I defended in Chapter 3: namely, that acts of attention are subjectively inalienable.

4.2 The sharpening and boosting of perceptual appearances

This section introduces some of the empirically documented effects that covert visual attention has been shown to have on visual appearances. I will be focusing in particular on two such effects: what I will call the “sharpening” and “boosting” of visual appearances. While I do not believe these effects *exhaust* the contribution that attention makes to visual appearances (for reasons I’ll explain in §4.4), they are all that we need to grasp the replication argument and the points that I wish to raise in response to this argument. Having introduced sharpening and boosting as my paradigm cases of attention-induced appearance enhancements, I then outline the thesis that I will be defending in subsequent sections: the attentional appearance view.

At least on the face of it, conscious attention possesses its own phenomenology: there is “something that it’s like for you” to consciously attend (Nagel 1974). To get an initial grip on this phenomenology, it helps to consider the way conscious attention is described, both within folk psychological and empirical contexts. For example, conscious attention has been described as a “clear and vivid . . . focalization, concentration, of consciousness” (James 1890/1950); a “searchlight”, “spotlight”, or “zoom lens” (see Styles 2006, 65 ff); “experiential highlighting” (Campbell 2002); “phenomenal salience” (Wu 2011c),¹ and “prominence” (Ganson & Bronner 2013). It is also commonly described using spatial terms—e.g., as imposing “center-periphery” or “foreground-background” organization on experience (Watzl 2010; 2011; 2017).

Confronted with these descriptions, an initially attractive hypothesis is that each of

¹In this context, “salience” does not have the same meaning it usually has in the empirical literature. In the empirical literature, “salience” typically refers to the property of being attention-capturing or possibly the intrinsic basis of that dispositional property (see §2.3). In the present context, “salience” refers instead to the characteristic phenomenological upshot of conscious attention.

the above expressions picks out features that accrue to visual phenomenology in virtue of foveating an object—i.e., physically positioning the eye so that a target stimulates the foveal region of the retina, which is the area of the retina with greatest visual acuity and chromatic colour sensitivity. When one looks at an object overtly, for example, one is likely to experience it more *clearly, distinctly, and centrally* than when one experiences it out of the corner of one's eye—i.e., parafoveally. According to what we could call “the foveation view” of attentional phenomenology, the phenomenal contribution that visual attention makes to visual phenomenology reduces to the effects of foveation.

While there are ongoing empirical controversies about the explanatory relation between visual attention and oculomotor systems (Armstrong 2011; Smith & Schenk 2012), the possibility of covert conscious visual attention seems to show that we cannot simply *identify* visual attention's contribution to visual phenomenology with the effects of foveation (Posner 1980). To illustrate, pick an object in your immediate surroundings. Now, without taking your eyes off that object, shift your visual attention from it to something else in your surroundings. Most will agree that when they do this, there is an introspectively accessible phenomenological change accompanying the shift of attention from the object of fixation to the object spatially peripheral to the fixated object. Introspectively, what it's like to visually experience the object in the spatial periphery differs depending on whether or not one is focusing one's attention on it, even when foveation remains unchanged. There remains a sense in which the attended object feels clearer, more vivid, and more central in one's experience than the unattended one, even while the attended item is in the spatial periphery. This seems to suggest, against the foveation view, that the phenomenology of visual attention does not simply reduce to the effect of foveation on visual phenomenology. It seems to show a sense in which objects can be in some sense subjectively “central” within one's experience despite not being central to foveation.

Although the contribution that conscious attention makes to visual appearances cannot be reduced to the effects of foveation (contra the foveation view), it may be that acts of attention nevertheless *mimic* certain effects that foveating has on visual phenomenology. This is approximately Stazicker's (2011a; 2011b) view of (some instances of) conscious covert visual attention. Building on the empirical results of Yeshurn and Carrasco (1998), Stazicker argues that that “attending enhances spatial resolution in the visual signal itself, by mimicking the effect of foveation, by effectively making the spatial filters for a location more finely tuned”—an effect sometimes called “position” or “filter tuning” (Stazicker 2011b, 180). As Stazicker interprets these findings, some shifts of conscious covert visual attention consist in increasing the determinacy with which visual experience represents properties. By this, he means that attending makes it the case that a visual experience represents more determinate or precise properties than would be represented outside the

focus of attention, thereby enabling the subject to better resolve fine details of information at the attended location. I'll call putative cases in which attending achieves enhanced visual determinacy (e.g., increased spatial resolution via filter tuning) cases of "sharpening".²

Apart from sharpening visual experience, attending has also been found to boost the apparent degree of magnitude of various visual parameters. To illustrate a core finding, consider the following experiment from Carrasco et al. (2004) investigating the effect of visual attention on the appearance of contrast. In their experiment, subjects are told that they will be shown two "Gabor patches" (one on the left and one the right) (see Figure 1 for examples of Gabor patches with varying levels of contrast). Subjects are given the task of reporting the line orientation of the Gabor patch that looks higher in contrast ("is the stimulus that looks higher in contrast tilted to the left or to the right?"). In one condition they see a neutral cue at the fixation point, and in another condition their attention is cued (automatically and covertly) to the left or right side of a visual display by a direct cue. In each case, the Gabor patches appears about 50 milliseconds after the cue offset, and is visible for about 40 ms. Subjects have 1 second to respond. The basic result is this: if two Gabor patches of equal contrast are shown (one on the left and one on the right), subjects report the patch at the cued location as looking higher in contrast than the other patch. Similarly, if one of the patches is a bit lower in contrast than the other, cuing attention to the location of the lower contrast patch will make it appear equal in contrast to the higher contrast patch. The difference to apparent contrast can be up to 6% when higher contrast patches are being used. For example, a patch with 22% contrast might be made to look equal in contrast to a patch with 28% contrast.

²See also Nanay (2010) for the proposal that visual attention contributes to perceptual experience by increasing representational determinacy. Like Stazicker, Nanay explicates representational determinacy in terms of the determinable-determinate relation. He suggests that a shift of visual attention consists in coming to visually represent a more determinate property than would be visually represented without attention. Whereas Nanay seems to put this proposal forward as a fully general account of the difference attention makes to perceptual phenomenology, Stazicker puts it forward as an account of some instances of conscious attention.




Figure 1: Reprinted by permission from Springer Nature and Copyright Clearance Center: Springer Nature: Nature Neuroscience: “Attention alters appearance” Carrasco, Ling, and Read (2004).

Carrasco and colleagues have argued that the effect is genuinely perceptual, rather than a consequence of post-perceptual decision bias or response bias.³ They explain the increase in apparent contrast in terms of an enhancement made to the visual signal—specifically, to an increase in the contrast gain of the visual representation of the stimulus. According to this explanation, when attention to a location is increased, the signal strength of a neuron responding to that location becomes amplified, mimicking the effect that an actual increase in stimulus contrast would have under conditions of inattention. This process of gain modulation is thought to work in concert with other mechanisms to improve signal-to-noise ratio and make attended stimuli more discriminable to the subject (Carrasco 2011, 1489). In this way, an intelligible relationship is found between

³Various critiques have been made to Carrasco’s methodology, including that the result might be due to decision bias (e.g. Schneider and Komlos 2011; see also Beck and Schneider 2017 for a more theoretical discussion). Carrasco and coworkers have responded to these methodological criticisms (e.g., Anton-Erxleben et al. 2011), and I will be following their interpretation here.

familiar neural signatures of visuospatial attention at early stages of visual processing (e.g., contrast gain modulation) and the behavioural data showing attentional changes in visual appearance. As Carrasco and colleagues put it, there is “a strong link between the attentional mechanisms improving visual processing and the concomitant effects on appearance” (Barbot et al. 2018, 1831). The suggestion, in particular, appears to be that the same properties of neural states—e.g., contrast gain modulation; filter-tuning, etc.—that explain enhanced visual processing (signal enhancement) also help to explain the enhanced appearance that subjects report. The former visual signal enhancements (e.g., amplified contrast representation) are reflected in the visual phenomenology.

Similar subjective “boosts” in degree of magnitude have been discovered for other visual parameters, including: gap size, speed, brightness, and saturation (though not hue) (see Wu 2014b, Ch. 4 for a helpful review).⁴ Because these attentional effects on visual appearances mimic an actual boost in the magnitude of a visual parameter (e.g., contrast), I shall call such effects cases of “appearance boosting”. Whereas visual sharpening plausibly connects with our intuitive idea of attention as a form of *focus*, visual boosting connects with our intuitive idea of attention as a kind of *highlighting* or *spotlight* that makes attended items appear more distinct from their surroundings. As Carrasco (2011) notes, it is likely that mechanisms of filter tuning and gain modulation work cooperatively to improve signal discriminability. Each provides a likely basis for the subjective changes in appearance that commonly accompany shifts of visual attention.

I’ve discussed two ways that reorienting visual attention (in this case, visuospatial attention) can make a difference to visual phenomenology: namely, to sharpen and to boost the appearance of properties like spatial resolution and contrast. Various philosophical controversies have arisen around these attentional effects on visual appearance. One such controversy is about which metaphysical theory of perceptual experience is best able to accommodate the observed effects of attention on visual appearances: whether “qualia” or “mental paint” theory (Block 2010; 2015); naïve realism (Brewer 2013); or “representationalism” or “intentionalism” (Stazicker 2011a; Watzl forthcoming). I want to set this particular controversy aside here and focus instead on a different one. The controversy I am interested in concerns what significance effects like sharpening and boosting have for our understanding of conscious attention and its phenomenology. I’ll conclude this section by introducing the controversy I will be considering.

Let the “attentional appearance view” be the view that: (i) a state of conscious visual attention consists in a certain enhanced visual state; and (ii) the phenomenal character of a

⁴A possible explanation for why attention does not alter hue, unlike saturation and the other tested visual dimensions, is that hue cannot in any obvious sense be “increased” or “boosted” in order to facilitate stimulus discrimination (Fuller & Carrasco 2006: 4043). A stimulus could be experienced to change in hue (e.g., from blue to purple), but this would not make it appear to have more hue.

state of conscious visual attention consists in the enhanced way things appear to the subject in having that visual experience. According to the attentional appearance view, a state of conscious visual attention to something consists in a certain *enhanced visual awareness* of the thing. So, for example, earlier when you turned your attention covertly from one object in your environment to another one, the status of the latter as subjectively “central” or “foreground” within your experience, according to the attentional appearance view, consists in its possessing a certain enhanced appearance. Opponents of the attentional appearance view deny that an enhanced visual appearance itself ever *constitutes* a state of conscious attention to an object. On this view, whatever enhanced visual appearances may be associated with conscious attention are merely the *effects* of one’s conscious attention *on* something non-attentional. On this view, we fail to identify the distinctive phenomenology that subjects access when introspecting an attentional episode—its characteristic “foreground-background” subjective organization—by citing an enhanced visual appearance of the attended entity. On this view, attentional phenomenology, itself, consists in something over and above any such changes in visual appearance.

An attentional appearance view can, in turn, be either reductive or non-reductive. According to a non-reductive version of the view, we explain the sense in which an object’s appearance becomes “enhanced” through attention in terms of its coming to look *more attended* or (as this term is sometimes used) more “salient” (see Chalmers 2004 for this suggestion). This view does not attempt to elucidate the meaning of “enhanced” in “enhanced appearance” in non-attentional terms. On a reductive version of the view, by contrast, “enhanced appearances” are appearances of a kind that are, in principle, also available without attention: for example, sharpened spatial resolution and boosted apparent contrast. In this chapter, I am defending a reductive version of the attentional appearance view. I will focus on the sharpening of spatial resolution and boosting of apparent contrast as my main examples of how an act of conscious attention enhances visual appearances. To be clear, my claim is not that sharpening and boosting exhaust the phenomenology of conscious attention. This is for two reasons. First, as I’ll discuss in §4.4, I doubt that effects like sharpening and boosting exhaust the contribution that conscious attention makes to visual appearances. But they are all that I need in order to make the points that I want to make in the following section. Second, I endorse the attentional appearance view as a thesis about states of conscious attention and specifically about their characteristic subjective organization. As we’ll see, it is important, in my view, that a state of conscious attention always manifests within a larger act of attention. I therefore reject the attentional appearance view as a thesis about total attentional phenomenology. Nevertheless, I think the view, appropriately qualified, is more secure than

recent philosophical work on attentional phenomenology has made it out to be.⁵

4.3 The replication argument

My aims in this section are the following. First, in §4.3.1, I outline the replication argument against the attentional appearance view. According to this argument, there is introspective evidence against the attentional appearance view in the form of a phenomenal contrast. Second, I argue that the replication argument is unsuccessful because, while I acknowledge the phenomenal contrast, there is an alternative explanation for the contrast that proponents of the argument have not considered, and moreover this alternative explanation is consistent with the attentional appearance view. Then, in §4.3.2, I argue that the neglected alternative is superior to the explanations that proponents of the replication argument have offered. If this is right, then not only is the introspective evidence put forward by the replication argument consistent with the view that it claims to undermine, but better explained by it. Appreciation of the flaws in the replication argument will, in turn, reveal the best version of the attentional appearance view.

4.3.1 Replying to the replication argument: Step 1

The replication argument against the attentional appearance view is developed most systematically by Watzl (2017, Ch. 9) and is also a central argument in Wu (2011c; 2014b, Ch. 4).⁶ The argument purports to show that the distinctive phenomenology that we access when introspecting conscious visual attention cannot consist in any property of how things look or appear to the subject. It purports to demonstrate this by constructing pairs of experiences that are phenomenally alike in the appearances they present, but phenomenally different overall because of how the subject's attention is distributed. The argument exploits the assumption that for any visual appearance that might accompany an attentional episode, this appearance can be replicated by judiciously altering properties of the stimulus. For example, to replicate the sharpening of apparent determinacy that occurs with attention, one might insert a pane of glass possessing differential resolution

⁵It might be noted that, among philosophers who have directly spoken to the issue, the attentional appearance view seems unfashionable. My focus in this chapter will be on Watzl and Wu because they have been the most explicit critics of the view. However, they are not alone in rejecting the claim that a state of conscious attention consists in a distinctive modification of visual appearances. Others who I read as rejecting that claim include: Speaks (2010); Jennings (2015); Pautz (2010); and Beck and Schneider (2017). Among these authors, Speaks, Jennings, and Beck and Schneider seem to endorse primitivist positions about attentional phenomenology. Pautz explores the view that conscious attention is a form of conscious thought, and Stazicker (2011a) defends this cognitivist view in the case of "voluntary" attention (though he defends a version of the attentional appearance view for the case of involuntary attention).

⁶Watzl, in turn, traces the argument to Husserl (1983/1913) (see Watzl 2017, 181-2).

and blurriness between the subject and the object (with higher resolution at the location of the relevant object). To replicate the boosting effect, one could increase the physical contrast of the stimulus. (Recall that the underlying neural process of gain modulation is thought to imitate the effect that an actual increase in stimulus contrast would have on the visual signal under conditions of inattention). And yet, on the face of it, a scenario in which the sharpened and boosted appearance is the product of attending (the “attention scenario”) differs phenomenally from a scenario where the same appearances are achieved as a result of environmental manipulation (the “appearance replica” scenario).

I will grant the claim that it is possible to replicate, without attention, every difference in visual appearance that occurs as a result of an act of conscious attention. Moreover, I will grant that the phenomenal contrast between an attention scenario and its corresponding appearance replica establishes that an enhanced (e.g., sharpened and boosted) visual appearance, on its own, is insufficient for the phenomenology of attention. I want instead to challenge two further conclusions that Watzl and Wu each draw from the possibility of appearance replicas. The problematic inferences are the following. First, they infer that the enhanced visual appearance of an attended object could not constitute an instance of conscious attention to the object but could at most be an *effect* of conscious attention *on* visual appearances. Second, they infer that the phenomenology of conscious attention consists in something other than an enhanced visual appearance.⁷ Though their positive accounts of attentional phenomenology are quite different, Watzl and Wu each take the possibility of appearance replicas to show that the subjective structure of conscious attention consists in something other than how things appear to the subject.

However, the possibility of appearance replicas offers little support to the conclusions Watzl and Wu draw from it. Appreciation of why this is so will point us toward the position that I will be defending in this chapter. I will focus on *shifts* of attention to make my point. In such a case, the corresponding appearance replica would be a scenario in which an object’s visual appearance changes because of changes occurring in the environment. For example, the object might come to appear higher in contrast because the object’s con-

⁷For example, after discussing the possibility of appearance replicas, Watzl concludes that any difference conscious attention makes to how things appear perceptually to the subject “do not flow from the nature of attention. While they are [causally] explained by attention, they are not constituted by attention. For this reason it is possible to have these effects also without attention” (2017,174). Watzl is explicit that “The reason we could find a replica for each effect of attention on the appearances was that all of these are *mere effects* or *correlates* of a certain distribution of attention”, rather than *constitutive* of attention (2017, 186, italics added). Similarly, when examining the empirical findings showing the many ways attention alters how things look to the subject, Wu writes that “This work in fact does not reveal a distinctive phenomenology of visual attention, but rather the effects of attention on visual phenomenology” (2014, 111) and “It is not clear . . . why one should speak of the resulting phenomenology as specifically attentional since it involves changes in *visual* phenomenology such as alterations in apparent contrast, size, and saturation. Couldn’t such changes occur in visual experience without attention?” (2014, 122, italics original).

trast increases and the object might come to appear clearer because the resolution of the intermediate pane of glass differentially changes.

Let's begin with Watzl's and Wu's claim that the enhanced appearance of the attended object does not constitute an instance of visual attention but is only an effect correlating with attention. Against this, one could hold that the enhanced appearance of an object constitutes conscious attention to the object only if the enhanced appearance is appropriately generated. In particular, it could be that an object's enhanced visual appearance to a subject constitutes the subject's state of visual attention to the object only if the enhanced appearance is generated endogenously by the agent's motivational states. This condition is not met in the appearance replica scenario where one tampers with the stimulus. This would make the case analogous to how, on causal theories of action, bodily movement qualifies as action only if it is internally generated by the agent's intention; otherwise, a qualitatively similar bodily movement would be merely a bodily movement. In the same way that a bodily movement is insufficient for an action, the enhanced appearance in an appearance replica scenario may be insufficient for attention.

For parallel reasons, the possibility of appearance replicas does not establish that the distinctive "foreground-background" or "center-periphery" structure that subjects access when introspecting an attentional episode consists in something other than a property of visual appearances. In particular, we can consistently reject the latter claim while acknowledging an additional phenomenal factor present in the attentional scenario that is absent in a corresponding appearance replica scenario. We can claim that whereas in the attentional scenario the enhanced appearance is experienced as having its apparent source in the agent, the same enhanced appearance in an appearance replica does not manifest as originating in the agent. Instead, it manifests as having an external origin (e.g., in the increasing resolution of the mediating pane of glass and in the physical contrast of what is behind it). It may be that the enhanced appearance of an object partly constitutes the phenomenology of attending to it (in particular, the subjective structure of the attentional state), even if the former experience is insufficient for the latter. This would arguably make the phenomenology of attention analogous with experiences of acting. In particular, an experience of one's own body moving partly constitutes an experience of acting, even though the former experience is insufficient for the latter. In each case, it could be that what more is required (i.e., to experience the enhanced appearance as a state of attention; to experience one's bodily motion as an action) is to experience the change as a change of one's own making. We could then acknowledge that the same visual appearance that is achieved by attending can be achieved without attending, but deny that whatever subjective organization an act of conscious attention contributes to consciousness must be something over and above how things visually appear to the sub-

ject. For the latter to manifest as a state of attention, one could argue, it must manifest as coming from oneself.

The neglected alternatives to which I've been drawing our attention provide the core of the position that I want to defend. It also has clear echoes in a view advanced by James (1890/1950):

[I]t must be admitted that to some extent the relative intensity of two sensations may be changed when one of them is attended to and the other not ... But, on the other hand, the intensification which may be brought about seems never to lead the judgment astray. As we rightly perceive and name the same color under various lights, the same sound at various distances; so we seem to make an analogous sort of allowance for the varying amounts of attention with which objects are viewed; and whatever changes of feeling the attention may bring we charge, as it were, to the attention's account, and still perceive and conceive the object as the same. (James 1890/1950, 426)

The explanation that I'm outlining of the phenomenal contrast between an attention scenario and an appearance replica makes a point similar to the one that James makes in this passage. My explanation claims that changes in attention consist in changes in the subjective intensity (as well as "clarity", as James goes on to discuss) of visual appearances which we experience as of our own making rather than as originating in the world. If we combine this with James' thesis that attention is a kind of "focalization" or "concentration" of consciousness ("in clear and vivid form"), the way is open to understanding the phenomenology of conscious visual attention as a *subjectively owned perceptual enhancement*. Similar to how adopting a different spatial perspective on an object allows you to vary how the object appears to you without it appearing to change, you can vary how an object looks to you by moving your conscious attention around without the object (or anything else in the environment) appearing to change.

So far, I've been assuming scenarios in which an object's properties remain constant throughout an attention shift. My proposal has been that, in the attention scenario, while the object's properties remain unchanged, its appearance changes (e.g., sharpens and boosts) as a result of attention and this appearance change manifests as endogenously generated by the agent.

However, we can imagine scenarios in which the object's appearance changes at the moment the subject reorients her attention to it. For example, it could be arranged that as the subject's attention turns to the left, the physical contrast of the object on the left decreases and the resolution of the mediating pane of glass becomes a bit blurrier at its location. In this way, we could "cancel" whatever sharpening and boosting of appearance that the shift in attention would otherwise have caused. Does my proposal predict,

counterintuitively, that under these circumstances the agent would lack any experience of having shifted her attention, because the object's visual appearance remains constant throughout the attention shift? No. What matters for the experience of shifting attention to an object is not that the apparent contrast and determinacy of the object increase as one turns one's attention to it. What matters is that one experiences the object's appearance as partly depending on one's act of attention. In particular, one must experience it as possessing an enhanced appearance relative to the appearance that it would have were one's attention directed elsewhere. The comparison class relevant to determining whether a given appearance is "enhanced" is therefore not synchronic (since one might be attending to the lowest contrast or blurriest object in a scene), nor diachronic (since one could shift one's attention from something of higher apparent contrast object and resolution to something of lower apparent contrast and resolution). Rather, the comparison class is counterfactual: one experiences the object's current appearance as enhanced relative to the appearance that it would have were one not attending to it or were one's degree of attention to it different. This accommodates cases in which one experiences oneself attending despite the object of one's attention decreasing in apparent intensity. Even while its apparent contrast diminishes, one experiences its apparent contrast as higher than what it would be were one not attending to it as one is.⁸

I've suggested that the evidence that the replication argument invokes is consistent with the view that it claims to exclude. It remains to be seen whether there is, apart from consistency with the introspective evidence, any reason to prefer one explanation over the other. I'll now explain why we should prefer my version of the attentional appearance view over either of the rival explanations Watzl or Wu offer.

4.3.2 Replying to the replication argument: Step 2

I'll begin by arguing that there are undesirable empirical and theoretical consequences to claiming, as Watzl and Wu each do, that conscious visual attention is the cause of an enhanced visual appearance rather than partly constituted by the enhanced appearance. After that, I'll examine Watzl's and Wu's positive proposals for how we should accommodate the phenomenal contrast between an attentional scenario and its corresponding appearance replica. I'll suggest that each problematically divorces the subjective structure

⁸When considering the possibilities for a reductive proposal of the phenomenology of conscious visual attention in terms of apparent contrast magnitude, Wu (2011c; 2014b) considers proposals in terms of increased apparent contrast relative to other perceived objects at a time ("synchronic salience") and relative to the same object at different times ("diachronic salience"). Having argued that the phenomenology of conscious visual attention (or "phenomenal salience") does not track either of these properties, he concludes that phenomenal salience is not a property of visual phenomenology, but rather of cognitive phenomenology. He does not consider the proposal I've sketch here.

of conscious attention from the subject's perceptual perspective on the world. By contrast, my proposal doesn't have this result. So, we should prefer my proposal to theirs.

To begin with, the claim that conscious attention causes (rather than partly consists in) an enhanced appearance seems to bring with it undesirable empirical commitments. To see this, recall the suggestion from Barbot et al. of "a strong link" between properties of visual neurons that underlie the enhanced visual signal (e.g., filter-tuning; gain modulation) and the accompanying enhancements of visual appearance (e.g., sharpening and boosting) (2018, 1831). As I noted earlier, the link they seem to have in mind here is explanatory: namely, the same properties that explain visual processing enhancements also partly underlie the enhanced appearance subjects report having when they are attending. For example, we may experience the attended object as appearing higher in contrast than when it is unattended because the gain of the visual signal is amplified in the attended case relative to the unattended case. The amplified gain, the thought is, explains both improved stimulus discriminability and its boosted appearance. So, when opponents of the attentional appearance view assert that attention is the cause of changes in an object's visual appearance, either they are affirming that attention is also the *cause* of the processing enhancements that are realized by gain modulation and filter-tuning or they are rejecting the explanatory connection that these researchers are drawing between signal enhancement and enhanced visual appearance. The latter would be undesirable, given that this is a leading account in the area. So, let's consider the former claim—namely, that visual attention is not itself the visual signal enhancement, but is rather the causal source *driving* the signal enhancement.

In terms of the source-resultant framework that I laid out in Chapter 1, the natural way to read the above suggestion is as suggesting that attention is realized by the "supramodal" source of bias rather than the resultant effect of biasing on visual processing. However, this would be at odds with an emerging empirical consensus that takes attention not to be the top-down source of bias into the visual system, but the globally integrated state into which the visual system settles as a product of biased competition (Desimone and Duncan 1995; Duncan 1998).⁹ I do not think that either Wu or Watzl wish to identify attention with the source state of attentional bias. Wu explicitly argues

⁹Admittedly, some researchers working within the biased competition framework write as though they maintain allegiance to the more traditional conception of attention as the higher level cause of signal enhancement (Kastner 2009; see Mole 2015 for discussion). But even these researchers can be interpreted as claiming that there is a dedicated set of source regions for attentional *control*. The question of whether or not there are dedicated mechanisms responsible for attentional control (e.g., within frontal-parietal areas) seems to me an open question that can coherently be posed within the biased competition framework without abandoning commitment to a conception of attention as the resultant state of biased competition. The view that there is a privileged set of attentional control mechanisms does not, by itself, signal a return to the view of attention as the cause of selectively enhanced processing.

against the supramodal conception on theoretical and empirical grounds (see especially Wu 2017, 24-5), and Watzl makes empirical claims that contradict it as well—e.g., by relating the neural correlates of visual attention to visual areas (e.g., 2011, 180) and relating the driving states of visual attention (in both “active” and “passive” cases) to the LIP priority map (see Watzl 2017, 132, 147). Perhaps Watzl and Wu will reply that attention is best understood as functionally mediating between the source of attention and the resultant state of visual signal enhancement. But while there could be intermediate stages in the attentional process, it is unclear what would be gained theoretically by denying that the enhanced visual signal itself exemplifies a form of attention once one has abandoned the supramodal conception of attention. It is simpler to claim that the process of visual attention *just is* the biasing of visual states by source states, and that a state of visual attention *just is* the enhanced visual state that results from that process (as realized by, e.g., properties like filter-tuning and gain modulation) (cf. §1.3). The move to regarding attention as a further element *in addition to* the source state, the resultant state, and the bias process seems theoretically unmotivated.

One response we should guard against at this point would be the reintroduction of the replication argument at the level of processing enhancements. In particular, someone might propose the following argument. We cannot identify the neural basis of visual attention with states of selectively enhanced neural response within the visual system (e.g., gain modulation) because it is possible to duplicate the latter effects in non-attention involving ways. In principle, one could duplicate the same increase in signal strength that accompanies a shift of visual attention simply by strengthening the physical stimulus, so that the resulting visual signal is qualitatively identical to the one that would have resulted from a weaker stimulus under attention. Consequently, the neural basis of visual attention cannot *consist* in an enhanced visual signal. Rather, it must be the *cause* of that signal enhancement. In response, though, this argument is no more persuasive at the physical level of neural realizers than it is at the phenomenological level of appearances. It could be that the enhanced visual signal qualifies as a basis of visual attention only when it is appropriately endogenously caused.

In summary, there are empirical consequences to claiming that the relation between attention and an enhanced visual appearance is causal rather than partly constitutive. In particular, given the likely explanatory links between enhanced visual appearances and enhanced visual processing, opponents of the attentional appearance view commit to a questionable view of attention as the causal force that drives selectively enhanced processing, rather than the achievement of states of selectively enhanced processing.

I will now turn to consider the rival explanations that Wu and Watzl give for the relevant phenomenal contrast, and I will suggest that my proposal is preferable. On the view

that I'm defending, recall, a state of visual attention is simply a motivationally modulated visual experience, and the distinctive subjective organization of conscious attention is a property of how things appear to the subject in this experience. On this view, the subjective organization of a conscious state characterizes the agent's *perceptual* perspective on the world. By contrast, when we examine the positive accounts that Watzl and Wu offer for the phenomenal contrast between an attentional scenario and an appearance replica scenario, we find each of them, in different ways, characterizing this subjective organization more meta-cognitively. Let's consider their proposals now.

The point is clearest in Wu's case. Wu proposes to explain the phenomenology of conscious attention in terms of a subject's introspective awareness. In particular, on Wu's view, the "foreground-background" structure that subjects access when introspecting attention belongs to the thought that one would express by saying "I am attending to *that*" (2011c, 95; 2014b, 130). On this view, the phenomenology that one accesses when introspecting an episode of attention turns out not to be a property of one's attention but of the *thought* that one is attending. As Wu points out, this implies that, outside of special contexts of introspection, visual attention does not intrinsically contribute to the subject's conscious perspective. If one is immersed in a game of squash or a climb, on his view, the foreground-background organization of attentional phenomenology is absent: nothing stands out as "central", "prominent", or "salient" to one in the way that it does when one consciously reflects on one's attentional state (2014b, 130). Wu appears to acknowledge that he is saying something surprising to common sense here. It is natural to understand him as attributing an introspective illusion regarding the ubiquity of attentional phenomenology—as though it shows up in experience only when we self-consciously reflect on attention. This is an unattractive result.

On the face of things, Watzl's account of the subjective structure of conscious attention does not have the same shortcoming as Wu's. Watzl's account does not tie attentional phenomenology to explicit self-reflection. On Watzl's view, conscious attention is the subject's mental activity of "structuring" her overall conscious state into what is more "central" and what is more "peripheral". This is an activity subjects engage in continuously as they act in the world, not only in special contexts of reflection. Taken on their own, Watzl's claims about conscious attention are intuitive and plausible. However, their plausibility diminishes on examination. We can bring this out by considering the difference between the type of conscious perspective subjects possess in virtue of acts of attentional structuring, on Watzl's view, and other examples of perspective. For example, whereas one's egocentric spatial perspective contributes to orienting one to the world around oneself, an act of conscious attention, on Watzl's account, orients one to one's *experiences* of the world. Through acts of conscious attention, experiences, not their objects, manifest

subjectively as central and peripheral to one another.

An apparent consequence of Watzl's account is that foveating an object and covertly attending to an object turn out to be acts with fundamentally different targets and aims. By foveating, one arguably aims to enhance one's *perceptual* relation to something in one's environment—e.g., by bringing a thing more clearly into view and into the center of one's visual field (with accompanying processing advantages). By contrast, on Watzl's account, conscious attention aims at enhancing the subject's relation to her own conscious states—e.g., by bringing an *experience* of a thing out from the conscious periphery and into the center of consciousness. If conscious attention endows subjects with a distinctive kind of perspective, then, on Watzl's view, it endows them with a perspective on experiences. Although Watzl does not wed the phenomenology of attention to explicit contexts of self-reflection like Wu does, his account nevertheless is like Wu's in being metacognitive. In this regard, Watzl's account, like Wu's, draws a sharp distinction between a subject's *perceptual* perspective on objects and her *conscious attentional* perspective.¹⁰

The natural question to ask of Watzl's account is whether acts of conscious visual attention really *are* so different from acts of foveation or other activities which aim at enhancing the subject's perceptual awareness of select parts of the environment. Of course, the two *are* different: foveation incorporates overt movement while covert attention does not. But beyond this, it is natural to suppose that they are activities with similar aims and targets. Each is a way of bringing something into *focus* and making it more visually discriminable from its surroundings. Each aims at seeing certain parts of the environment better. Covert visual attention may simply be one aspect of a more encompassing perceptual activity of "looking" around at the world: an activity that consists in actively

¹⁰Watzl employs perspectival metaphors in connection with attention, but it is sometimes unclear what these amount to on his view. Sometimes he characterizes conscious attention as itself a mental perspective (Watzl in preparation). When he does so, he is careful to note that whereas more familiar examples of perspective manifest as a perspective on *objects* (and thereby as structuring how the world perceptually appears to the subject), attention is a perspective on *appearances*. This gives the impression of a higher order view of attention as a perspective on conscious states ordered in terms of which experiences are more central to which others. Other times, Watzl avoids suggesting that attention is a separate type of perspective and instead describes it as simply the "structure" of one's conscious perspective on the world:

Attention structures consciousness into what is more central and what is more peripheral. Unlike spatial structure, the center-periphery structure of consciousness is not a structure of how the world appears to the subject through her conscious perspective. It is the structure of her perspective itself. (Watzl 2017, 183)

These descriptions remove the suggestion of conscious attention as higher-order, but they also risk collapsing Watzl's account into a version of the attentional appearance view. Suppose that the center-periphery structure of attention constitutes a subject's conscious perspective as a perspective on the world. In that case, we can ask what difference there is between the subjective structure of a perspective and the subjective structure that characterizes "how the world appears to the subject through her conscious perspective". Plausibly, it is the structure of a subject's conscious perspective that determines how things appear to a subject from that perspective. If so, then it is not clear there really are two "structures" in that case.

manipulating one's perceptive point of view in order to enrich one's information about a thing (cf. Matthen 2014, 2005). Indeed, some researchers argue that covert visual attention may simply be the visual processing enhancements that occur as a by-product of saccade planning in oculomotor areas, with covert attention possibly being the product of a sub-threshold motor command (see Armstrong 2011 for discussion). Such proposals are not obviously in conflict with the phenomenology of covert attention to an object (which can feel a bit like one is trying to see a thing better while simultaneously suppressing the movement of one's eye).

Watzl and Wu are led to their respective positions about conscious attention in light of the perceived deficiency of the claim that the phenomenology of conscious attention is a certain manner of enhanced visual appearance—e.g., a sharpened and boosted appearance of an object. That each takes the attentional appearance view as their starting point is an implicit indication that it is taken to be the default position. It is simple and it is naturally suggested by the empirical work on attentional alterations in visual appearance. In light of the perceived inadequacies of the position, these authors are led to characterize the phenomenology of conscious attention meta-cognitively. They are led to characterize it either as a subjective organization that is unique to contexts of introspective reflection or as the phenomenological manifestation of the order subjects give to their conscious mental states. I've argued that there is a simpler view available. Conscious attention is a motivationally biased conscious visual state. And the characteristic phenomenology that we access when we introspect conscious attention is simply the reflection of this motivational bias on the way things look to the subject, with some things being clearer and more distinct than others. Assuming that each account is phenomenologically adequate, mine is preferable. But some will object at this point that my proposal is *not* phenomenologically adequate. I'll turn to this objection now.

4.4 Modulating perceptual organization with attention

Opponents of the attentional appearance view might acknowledge that my proposal in terms of subjectively owned changes in visual appearance can explain *aspects* of attentional phenomenology. For example, they might acknowledge that my proposal can accommodate those aspects of attentional phenomenology that we intuitively describe using terms like “clarity”, “vividness”, and “distinctness”. Perhaps, they will grant, the aptness of such language does come from subjectively owned changes in an object's visual appearance like sharpening and boosting. Nevertheless, they may insist that my proposal falls far short of accommodating attentional phenomenology as a whole. This is because, they might suggest, my proposal cannot adequately accommodate the holistic

and structural aspects of conscious attention that give certain spatial metaphors their intuitive meaning. In particular, some may suggest that my proposal cannot accommodate those aspects of attentional phenomenology that motivate us to describe an attended entity as being *more central* or *more foreground* in our experience than unattended ones. And one might doubt that the latter aspects of attentional phenomenology *can* be accommodated exclusively in terms of a clearer, more distinct visual appearance, together with subjective ownership. To see this, consider that while shifting attention covertly to an entity might mimic some of the effects that occur as a result of foveation or environmental manipulation (e.g., making a thing appear literally clearer and more distinct from its surroundings), an attended entity need not *appear* literally more central to the subject than an unattended one. A covertly attended entity appears to you as being in the spatial *periphery*. And yet, when focusing one's attention on the entity, it *is* (in some possibly primitive sense) *more central* in your experience. Analogous points apply to the notion of "foreground", since what is in the attentional foreground can be in the spatial background. One might thus suggest that it is this holistic structure which agents actively endow on their experiences through acts of conscious attention. And one might further suggest that it is this aspect of attentional phenomenology that most obviously resists reductive characterization in terms of modulations in an object's visual appearance. If this is right, then conscious attention contributes something much more fundamental to experience than superficial changes like apparent contrast and spatial resolution. It contributes novel organization that is simply absent from the appearances themselves. This is the real reason, one might argue, the attentional appearance view is insufficient.

In response, proponents of the attentional appearance view should acknowledge that effects like sharpening and boosting are insufficient to accommodate attentional phenomenology. This is so even when sense of ownership is acknowledged to accompany these changes in appearance. However, they should reply that it would be a mistake to reject the attentional appearance view on these grounds. One can grant that effects like sharpening and boosting are insufficient to capture attentional phenomenology but argue that the former do not exhaust the explanatory resources available to an attentional appearance view. Even on the reductive version of the view that I am pursuing, one need not understand conscious attention's contribution to visual appearances only in terms of qualities like contrast and determinacy. There is, after all, more to a visual appearance than a bundle of qualities. Consequently, there are potentially more ways that an act of attention could alter a visual appearance than by altering the apparent magnitudes of those qualities. In particular, all will agree that visual appearances themselves possess complex internal structure. This includes principles of "gestalt" organization (see e.g. Wagemans et al. 2012). For example, visually resembling objects typically appear as be-

longing together as members of a single visual group, whereas visually dissimilar objects appear to stand apart from each other as members of different groups. Similarly, there are many well-understood principles governing how one entity comes to be separated or segmented from its surroundings. Beyond these gestalt principles of visual organization, appearances also instantiate a complex “bound” organization. For example, when seeing an object, one sees the thing as possessing multiple features—e.g., a certain colour and shape. Further, we do not typically see only one object in isolation, but multiple objects as spatially related to one another and to ourselves within one, overarching scene. These are all aspects of what I mean by the “internal organization” of a visual appearance.

Confronted with doubts regarding the adequacy of the attentional appearance view, a natural strategy for a proponent of this view is to ask whether they can invoke any of the above organizational features of visual appearances to alleviate such doubts. This would require asking: (a) whether acts of conscious attention might alter the internal organization of an appearance in something like the way that we have seen it to alter subjective magnitudes like contrast and spatial resolution; and, if so, (b) whether this, when combined with subjective ownership, could remove lingering doubts that drive theorists away from the attentional appearance view. This is the question to which I’ll now turn.

Let’s begin with whether attention may contribute to the internal organization of a visual appearance. While the topic of attention and its relationship to visual organization is a complex one with a long history, current work suggests that visual attention is a modulator of perceptual organization and that the relationship between the two is “multifaceted and mutually constrained” (Kimchi et al. 2016, 34-5). For example, de Haan and Roden (2010) investigated whether similarity grouping—a principle of perceptual organization traditionally assumed to be established pre-attentively—is modulated by “attentional relevance”. They gave subjects a task involving the simultaneous identification of two targets (one on the left side of the display and one on the right side). They found that if the two targets are similar along a task-relevant dimension (e.g., visually similar in letter identity in a letter identification task), performance improves. This much, they note, is consistent with the received assumption that the competitive interactions that culminate in a state of visual attention occur among already organized groups of partially bound visual individuals and, furthermore, that elements that are grouped tend to cooperate rather than compete for processing. This is consistent with a picture on which the process of biased competition that culminates in a state of visual attention operates over a visual field already organized along gestalt principles. But de Haan and Roden found that when the two targets are similar along a task-irrelevant dimension (e.g., in colour identity in a letter identification task), similarity grouping was not observed. The

result, they note, is surprising (though, they also note, consistent with recent work): “supposedly pre-attentive grouping mechanisms might not operate as independently from top-down attentional modulations as traditionally thought” (de Haan and Roden 2010). Similar results have been found with respect to other gestalt principles of visual organization. For example, whether an ambiguous visual figure is treated by the visual system as a dumb-bell or as a circular individual within in a larger dumb-bell—as measured by the “same object advantage”—has been found to vary depending on whether the agent’s intention is to grasp (in which case attention spreads to throughout the dumb-bell form) or point (in which case attention spreads only through the embedded circle) (Riddoch and Humphreys 2007).¹¹ It seems that attention neither determines nor is determined by a visual state’s internal gestalt organization. Rather, attending *modulates* such organization.¹²

The above work on the modulatory effects of attention of perceptual organization does not speak directly to the question of visual appearances. Less empirical work has been devoted to the modulatory effects of attending on the appearance of visual organization than to lower level properties like resolution and contrast magnitude (though Carasco and colleagues have recently begun to investigate these questions see Barbot et al. 2018). But we perhaps do not need to wait on empirical evidence. The claim that attention alters visual organization is a familiar theme from the Gestalt psychological tradition and familiar introspectively from examples of ambiguous figures and examples of aspect (e.g., the duck-rabbit; the Necker cube) (Wagemans et al. 2012). It may also be a central part of the experience of searching for something. When you look for a thing in a crowded scene, you are intuitively attempting to “single it out” from its surroundings. This will be easy if pre-attentive visual organization ensures that your target stands apart from its surroundings—e.g., when there is weak target-distractor grouping (Duncan and Humphreys 1992). But it will be quite difficult when your target is perceptually grouped with its surroundings in virtue of visual resemblances—i.e., when there is strong target-distractor perceptual grouping. Because competitive interactions tend to occur between rather than within perceptual groups, all members of the same perceptual group will tend

¹¹The “same object benefit” refers to the finding that if a subject’s target appears inside the same enclosed figure as a previously presented cue, then the target is detected more quickly than if the target appears equidistant from the cue but located within a different enclosed figure (Egley et al. 1994). There are controversies surrounding the interpretation of this and related findings, though it is commonly cited as evidence that attention operates upon already segmented visual objects, rather than being, for example, location-based.

¹²A similar conclusion seems to be true for feature-binding. In a review article, Humphreys (2016) concludes that “rather than thinking of visual binding being either fully bottom-up and preattentive, or fully dependent on attention, it may be better to think of bottom-up and top-down processes interacting to optimize feature binding” (2016, 1933). He also puts the point by saying that it is more accurate to regard attention as “modulating” binding than as being either necessary for it or irrelevant to it.

to be selected together and this will frustrate your goal of identifying a specific individual within that group.¹³ In searching for the target, perhaps part of what you are attempting to do is reorganize a scene's appearance to align it with your goals: to pull the target out from the crowd to which it has been pre-attentively assimilated by similarity grouping. You are trying to modulate the gestalt organization of the appearance so as to foreground or single out your target from the crowd. This would be an active, intention-guided process of perceptual reorganization.

Once we expand our view of the potential ways conscious attention can impact perceptual appearances—including not only lower level changes like the sharpening of spatial resolution and the boosting of apparent contrast but also gestalt principles of perceptual organization—it is unclear why we should conclude that the attentional appearance view is phenomenologically inadequate. In particular, it is not clear that the features of attentional phenomenology that motivate describing conscious attention as form of “structuring” of experience into relations of “foreground” and “background” or “center” and “periphery” are necessarily left out of the attentional appearance view (though specific proposals would need to be put forward and assessed). We must, in particular, be careful to distinguish the experience of partial ownership over the internal organization of a visual appearance and the experience of a novel organization that originates with acts of conscious attention and that is intrinsically foreign to appearances. The relative merits of these proposals should not be settled solely on the basis of introspection. We must appeal to broader theoretical considerations of the kind I appealed to in the previous section. Until each hypothesis is properly examined, it is premature to abandon the attentional appearance view on the grounds that it cannot accommodate the holistic or gestalt organization of conscious attention. Such features of attentional phenomenology may be amenable to the same treatment that I earlier proposed for sharpening and boosting.

4.5 What appearance replicas show: the inalienability of conscious attention

Recent philosophical work on the nature and phenomenology of conscious attention has emphasized the theoretical significance of the possibility of appearance replicas. The driving intuition is that experienced changes in an object's visual appearance are insufficient to accommodate what it's like for subjects to attend consciously to an object. This insufficiency is thought to be revealing both about the nature of conscious attention and

¹³This is also one account of the “flanker interference effect” discussed in Chapter 1 (Driver and Baylis 1989; see §1.5)

how it figures in total phenomenology. In particular, some have cited the possibility of phenomenally contrastive appearance replicas as evidence that the distinctive subjective organization of conscious attention consists in something over and above any organization within appearances. This position acknowledges that conscious attention has various effects on how things look and it acknowledges that the “look” of things has its own, highly complex internal structure. But it insists that acts of conscious attention add their own structure to experience. This further structure reflects, to use Lewis’s phrase, an agent’s “constructive acts” through which the agent confers selective emphasis on a passive appearance.

I’ve argued that these are the wrong conclusions to draw from reflection on phenomenally contrastive appearance replicas. Conscious attention does not contribute an additional structure to experience. Rather, the subjective structure of conscious attention is a structure within appearances. At the same time, appearances are not passively given or presented to the subject. Instead, appearances are shaped under the guidance of our priorities. This activity is reflected in the phenomenology—e.g., in the experience of aspects of an appearance as the product of our own activity.

I want to close by suggesting that we are in a position to draw a different conclusion about the nature of conscious attention from the possibility of phenomenally contrastive appearance replicas. This conclusion relates to an implication of the account that I defended in Chapter 3—namely, that acts of conscious attention are subjectively inalienable. Recall that, on the view defended there, an agent subjectively owns any action that is controlled by a motivational state that controls conscious attention. This implies that if conscious attention is an action, then, necessarily, it is a subjectively owned action, since whatever motive controls the act of conscious attention is (trivially) controlling conscious attention. Put differently, it is impossible to experience an episode of conscious attention as an act of conscious attention, but as one that is not performed by oneself. Acts of conscious attention are subjectively inalienable. It follows that there can be no analog in the domain of conscious attention of schizophrenic passivity experience (see §3.6).

This impossibility claim might give us pause for thought. Specifically, one might wonder whether there are any reasons independent of my account for accepting the impossibility claim.¹⁴ I’ll close by suggesting that the explanation that I’ve defended in this chapter for the phenomenal contrast between an experience of conscious attention and its corresponding appearance replica gives us independent support for this impossibility claim. In so doing, I will be using the results of this chapter to reinforce the account of subjective action ownership defended in Chapter 3. I will focus my remarks on subjective ownership for shifts of conscious attention, though parallel remarks can be made about

¹⁴Thank you to Christopher Mole for pressing this worry to me in an earlier commentary on my material.

maintaining states of conscious attention.

First, let's recall examples of actions that appear to be only contingently subjectively owned—e.g., an alienation experience like anarchic hand or schizophrenic alien control. If we take patients' subjective reports of these experiences at face value, they seem to show that it is possible for a subject to lose subjective ownership over an action like reaching for something or moving one's fingers, yet remain aware of the movement as an action or intentional movement. The subject remains aware of a certain action as being performed, but experiences herself as not performing the action. For example, a patient might report feeling as if an alien agent had entered her body and guided its movements (Spence et al. 1998, 1998). This seems to describe a kind of experience of action, but one that lacks the feature of first-person ownership or "mineness" that characteristically accompanies action. It is a pathologically third-personal experience of acting. It may be like observing another agent act (except, strangely, with one's own body).

In order for a shift of conscious attention to be subjectively alienated, the following would need to be the case. The subject remains aware of an attention shift as an attention shift, but experiences this shift as not her own doing. It is the possibility of such an experience, I want to suggest, that reflection on the possibility of phenomenally contrastive appearance replicas seems to foreclose.

We can develop these ideas more precisely as follows.¹⁵

1. An experienced change in an object's visual appearance is experienced either as a change in the object, a change in the object's surroundings, or as a change of the subject's own making.
2. If x is a shift in attention to an object, x is an experienced change in an object's visual appearance that is experienced as not being a change in the object or in the object's surroundings.

Therefore,

3. If x is a shift in conscious attention to an object, x is experienced as of the subject's own making. [From 1, 2].
4. In a case of subjectively alienated action, the subject experiences the action as performed, but as not performed by herself.

So,

5. There are cases of subjectively alienated shifts in conscious attention only if there are cases where the subject experiences a change as both of her own making and not of her own making. [From 3, 4]

¹⁵Thank you to Imogen Dickie for helping me with the details of this argument.

But,

6. There can be no such cases.

Therefore,

7. There cannot be cases of subjectively alienated shifts in conscious attention. [From 5, 6].

The argument is valid. So, let us assess the premises.

The first premise is controversial. This is because it is possible that an experience could remain “silent” about questions such as the causal source of an experienced change in an object’s appearance. For example, it could be that changes in visual appearance are neither presented as changes in the world (whether in the object or in its surroundings) nor as changes of one’s own making. Perhaps this is so in experiences like a fading afterimage after seeing a camera flash. We do not feel any sense of agency over these experienced changes of appearance, but arguably they also do not manifest as having occurred outside us.

In response, I mean to restrict the domain of objects at issue in the first premise to objects of perception that manifest as objective or independent of the subject. This includes experiences of what Siegel (2006) calls “object seeing”, which, as Siegel notes, differ experientially from sensation experiences like afterimages or phosphenes. One of the principal differences between experiences of object-seeing and sensation experiences, Siegel argues, is that the phenomenology of the former, but not the latter, reflects awareness of the subject’s perspective and capacities for interaction with the experienced object. For example, part of the phenomenal character of an experience of object-seeing is an awareness of how an object’s appearance will change systematically with changes in one’s perspective on it (“perspectival connectedness”). On this view, an experience of object-seeing constitutively involves experience of an object as being a certain way and of oneself as bearing a certain perspective on it. This view further suggests that the subject experiences an object’s appearance as determined by both the state of the object and by the subject’s current perspective on it. The first premise should therefore be interpreted as relative to the class of experiences Siegel calls “object-seeing”. It says, with respect to experiences of object-seeing, that an experienced change in an object’s appearance is either experienced as resulting from an environmental change—whether in the object itself or its surroundings—or as a change of one’s own making—i.e., as a result of a change in one’s perspective on it.

According to the second premise, a shift of conscious attention to an object is an experienced change in the object’s visual appearance that manifests neither as a change in the

object nor as a change in the object's environment. My support for this claim comes from the possibility of phenomenally contrastive appearance replicas. These cases suggest that when one experiences a change in visual appearance as having an outside source (e.g., whether as a result of a change in the object or in the mediating pane of glass), one is experiencing a fundamentally different sort of event than a shift of conscious attention. In particular, one is experiencing an appearance replica of a scenario in which one shifts conscious attention. By ensuring that a change in an object's appearance manifest to the subject as having a source outside the subject, we thereby remove the phenomenology of an attention shift. We do not preserve the experience of an attention shift but render it phenomenologically passive or alienated.

Opponents of the attentional appearance view will reject this appeal to appearance replicas. Specifically, they will reject the assumption that an experienced change in an object's appearance is a shift in conscious attention. On their view, these two things merely correlate. With the second premise, I am therefore presupposing the account that I have been defending in this chapter.

It follows from premises 1 and 2 that a shift of conscious attention is an experienced change in an object's appearance that manifests as of one's own making. More specifically, it manifests as a change in one's perspective on the object. This change in perspective can occur as the result of an intention—e.g., to improve one's perception of the thing, possibly in order to do something with the object. It can also occur automatically, as when one's attention is captured by a loud crash. In each case, to manifest as a change of attention, the change in appearance must manifest as arising from within the agent, rather than exogenously from the stimulus. Unlike other forms of perspectival change, there is nothing else that you need to *do* in order to change the object's visual appearance. For example, unlike changing your spatial perspective on an object, you do not need to first move your body in order to adjust your attention to the object. Nor do you need to do anything else (contra my opponents in this chapter). You change the object's appearance directly and covertly. That is what a shift of conscious visual attention consists in, on the present account.

Finally, claims 4–7 draw out the implications of the above observations for alienated experiences of attending: 4 states what I take an alienated experience of acting to consist in; 5 draws the inference that a subjectively alienated attention shift would have to be both an experience of oneself as both the agent of the shift and not the agent of the shift; and 6 claims that such cases cannot occur. Admittedly, it is difficult to conclusively demonstrate the last claim. However, we also do not have any clear reason to deny it. For example, when we examine actual examples of passivity experience, like those in anarchic hand syndrome and schizophrenia, we do not find subjective reports of being the agent of the

act and not being the agent of the act. The patients simply report feeling they are not the agent of the act. A plausible explanation for why such inconsistent experiences do not seem to occur is that an alienation experience simply consists in the absence of subjective ownership toward an action of which one is conscious. If an alienated experience of acting consists in an absence of subjective ownership over the action, then we have an explanation for why cases of simultaneous ownership and alienation cannot arise.

If this is right, then the possibility of appearance replicas does establish something important about attentional phenomenology. It establishes that it is impossible to experience a change in an object's appearance as a change in attention without thereby experiencing oneself as the source of the change. It shows that an act of conscious attention is essentially subjectively owned: an act of changing one's own perspective on an object.

4.6 Conclusion

In the last two chapters, I've defended a view of the phenomenology of conscious attention as the taking up of a motivational perspective. I've suggested that this element of perspective impacts our experience of ourselves as both agents and as perceivers. In an experience of acting, taking up a motivational perspective confers conscious ownership over an action, ensuring that it is the agent who controls what she does. In adopting such a perspective, the agent's perceptual experience of the world becomes subjectively organized according to the standards of success set by the motivational states that function to direct conscious attention. Consequently, the world perceptually appears to an agent in relation to her agential priorities.

However, we must treat the attentional structure of the agent's perceptual experience carefully. Specifically, I've argued that the case for primitivism about attentional structure is weaker than is sometimes acknowledged. The main evidence marshaled in its support—the possibility of phenomenally contrastive appearance replicas—does not provide clear support for the view. I've argued that the possibility of such cases is not only consistent with but better explained by a view on which attentional structure is the upshot of a modulatory process in which the agent experiences herself as an active participant. On the latter view, the agent does not, in adopting a motivational perspective, introduce a novel subjective structure to her experience that would otherwise be wholly absent. Rather, in attending consciously, the agent guides the construction of a perceptual appearance by biasing its development toward the enhancement of information that is relevant to her in light of her intentions and other priorities. This implies that the subjective structure of conscious attention is more interwoven with perceptual appearances than the primitivist can allow.

Bibliography

- Allport, A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action, in A. F. S. H. Heuer (ed.), *Perspectives on Perception and Action*, Hillsdale, New Jersey: Lawrence Erlbaum Associates, pp. 395–419.
- Allport, A. (2011). Attention and integration, in C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological Essays*, New York: Oxford University Press, pp. 24–59.
- Alvarez, M. and Hyman, J. (1998). Agents and their actions, *Philosophy* **73**(284): 219–45.
- Anderson, B. A., Laurent, P. A. and Yantis, S. (2011). Value-driven attentional capture, *Proceedings of the National Academy of Sciences of the United States of America* **108**(25): 10367–71.
- Anderson, B. A. and Yantis, S. (2013). Persistence of value-driven attentional capture, *Journal of Experimental Psychology: Human Perception and Performance* **39**(1): 6–9.
- Anscombe, G. (1957/2000). *Intention*, 2 edn, Cambridge, MA: Harvard University Press.
- Anton-Erxleben, K., Abrams, J. and Carrasco, M. (2011). Equality judgments cannot distinguish between attention effects on appearance and criterion: a reply to schneider, *Journal of Vision* **11**(13): 1–8.
- Armstrong, K. (2011). Covert spatial attention and saccade planning, in C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological Essays*, New York: Oxford University Press, pp. 78–96.
- Awh, E., Belopolsky, A. V. and Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy, *Trends in Cognitive Sciences* **16**(8): 437–43.
- Aydede, M. and Fulkerson, M. (2019). Reasons and theories of sensory affect, in D. Bain, M. Brady and J. Corns (eds), *The Philosophy of Pain: Unpleasantness, Emotion, and Deviance*, New York: Routledge, pp. 27–59.

- Baluch, F. and Itti, L. (2011). Mechanisms of top-down attention, *Trends in Neurosciences* **34**(4): 210–24.
- Barbot, A., Liu, S., Kimchi, R. and Carrasco, M. (2018). Attention enhances apparent perceptual organization, *Psychonomic Bulletin & Review* **25**: 1824–1832.
- Bayne, T. (2008). The phenomenology of agency, *Philosophy Compass* **3**(1): 182–202.
- Bayne, T. (2011). The sense of agency, in F. MacPherson (ed.), *The Senses*, New York: Oxford University Press.
- Bayne, T. and Levy, N. (2006). The feeling of doing: Deconstructing the phenomenology of agency, in N. Sebanz and W. Prinz (eds), *Disorders of Volition*, Cambridge: MIT Press, pp. 49–68.
- Beck, J. and Schneider, K. (2017). Attention and mental primer, *Mind & Language* **32**(4): 463–494.
- Beilock, S. L., Bertenthal, B. I., McCoy, A. M. and Carr, T. H. (2004). Haste does not always make waste: Expertise, direction of attention and speed versus accuracy in performing sensorimotor skills, *Psychonomic Bulletin and Review* **11**: 373–9.
- Belopolsky, A. V., Schreij, D. and Theeuwes, J. (2010). What is top-down about contingent capture?, *Attention, Perception, & Psychophysics* **72**(2): 326–41.
- Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience, *Physiology & Behaviour* **81**(2): 179–209.
- Bisley, J. W. and Goldberg, M. E. (2010). Attention, intention, and priority in the parietal lobe, *Annual Review of Neuroscience* **33**(1): 1–21.
- Blakemore, S. J., Smith, J., Steel, R., Johnstone, C. E. and Frith, C. D. (2000). The perception of self-produced sensory stimuli in patients with auditory hallucinations and passivity experiences: Evidence for a breakdown in self-monitoring, *Psychological Medicine* **30**(5): 1131–1139.
- Block, N. (1995). On a confusion about a function of consciousness, *The Behavioral and Brain Sciences* **18**: 227–247.
- Block, N. (2010). Attention and mental paint, *Philosophical Issues* **20**(1): 23–63.
- Block, N. (2015). The puzzle about perceptual precision, in T. Metzinger and J. Windt (eds), *Open MIND*, Frankfurt am Main: MIND Group.

- Bratman, M. (1996). Identification, decision, and treating as a reason, *Philosophical Topics* **24**(2): 1–18.
- Brent, M. (2017). Agent causation as a solution to the problem of action, *Canadian Journal of Philosophy* **47**(5): 656–73.
- Brewer, B. (2013). Attention and direct realism, *Analytic Philosophy* **54**(4): 421–35.
- Bridgeman, B. (1981). Segregation of cognitive and motor aspects of visual function using induced motion, *Perception and Psychophysics* **29**(4): 336–42.
- Briscoe, R. (2009). Egocentric spatial representation in action and perception, *Philosophy and Phenomenological Research* **79**(2): 423–60.
- Briscoe, R. and Schwenkler, J. (2015). Conscious vision in action, *Cognitive Science* pp. 1–33.
- Brogaard, B. (2011). Conscious vision for action versus unconscious vision for action?, *Cognitive Science* **35**: 1076–1104.
- Bruya, B. (ed.) (2010). *Effortless Attention: A New Perspective in the Cognitive Science of Attention and Action*, Cambridge MA: MIT Press.
- Bruya, B. and Tang, Y.-Y. (2018). Is attention really effort? revisiting daniel kahneman’s influential 1973 book attention and effort, *Frontiers in Psychology* **9**: 1–10.
- Buehler, D. (2014). *Psychological Agency: Guidance of Visual Attention*, PhD thesis, University of California, Los Angeles.
- Burnham, B. R. (2007). Displaywide visual features associated with a search display’s appearance can mediate attentional capture, *Psychonomic Bulletin & Review* **14**(3): 392–422.
- Byrne, A. and Eysenck, M. W. (1995). Trait anxiety, anxious mood, and threat detection, *Cognition and Emotion* **9**(6): 549–62.
- Campbell, J. (1994). *Past, Space, and Self*, Cambridge: The MIT Press.
- Campbell, J. (2002). *Reference and Consciousness*, New York: Clarendon Press.
- Carrasco, M. (2011). Visual attention: The past 25 years, *Vision Research* **51**(13): 1484–525.
- Carrasco, M., Ling, S. and Read, S. (2004). Attention alters appearance, *Nature Neuroscience* **7**(3): 308–313.

- Carruthers, P. (2015). *The Centered Mind: What the Science of Working Memory Shows us about the Nature of Human Thought*, Oxford University Press.
- Carter, J. D., Bizzell, J., Kim, C., Bellion, C., Carpenter, K. L., Dichter, G. and Belger, A. (2010). Attention deficits in schizophrenia: Preliminary evidence of dissociable transient and sustained deficits, *Schizophrenia Research* **122**(1-3): 104–12.
- Chalmers, D. (2004). The representational character of experience, in B. Leiter (ed.), *The Future for Philosophy*, New York: Oxford University Press, pp. 153–181.
- Chelazzi, L., Perlato, A., Santandrea, E. and Libera, C. D. (2013). Rewards teach visual selective attention, *Vision Research* **85**: 58–72.
- Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis, *Philosophical Transactions of the Royal Society* **362**: 1585–1599.
- Clark, A. (2001). Visual experience and motor action: Are the bonds too tight?, *The Philosophical Review* **110**(4): 495–519.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*, New York: Oxford University Press.
- Cochrane, T. (2018). *The emotional mind: A control theory of affective states*, Cambridge Cambridge University Press.
- Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain, *Nature Reviews* **3**(3): 201–215.
- Crane, T. (1988). The waterfall illusion, *Analysis* **48**(3): 142–147.
- Csikszentmihalyi, M. (1990). *Flow: The Psychology of Optimal Experience*, Harper Collins.
- Csikszentmihalyi, M., Abuhamdeh, S. and Nakamura, J. (2005). Flow, in A. Elliot and C. S. Dweck (eds), *Handbook of Competence and Motivation*, New York: The Guilford Press, pp. 598–608.
- Davidson, D. (1963). Actions, reasons, and causes, *The Journal of Philosophy* **60**(23): 685–700.
- Davidson, D. (1973). Freedom to act, in T. Honderich (ed.), *Essays on Freedom of Action*, Routledge.
- Davidson, D. (1978). Intending, *Philosophy of History and Action* **11**: 41–60.

- Day, K. A., Roemmich, R. T., Taylor, J. A. and Bastian, A. J. (2016). Visuomotor learning generalizes around the intended movement, *eNeuro* **3**(2): 1165–1182.
- de Haan, B. and Rorden, C. (2010). Similarity grouping and repetition blindness are both influenced by attention, *Frontiers in Human Neuroscience* **4**: 1–12.
- DeCaro, M. S., Thomas, R. D., Albert, N. B. and Beilock, S. L. (2011). Choking under pressure: Multiple routes to skill failure, *Journal of Experimental Psychology: General* **140**(3): 390–406.
- Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention, *Annual Review of Neuroscience* **18**: 193–222.
- Deubel, H. (2008). The time course of presaccadic attention shifts, *Psychological Research* **72**(6): 630–40.
- Deubel, H. (2014). Attention and action, in K. Nobre and S. Kastner (eds), *Oxford Handbook of Attention*, New York: Oxford University Press, pp. 865–292.
- Dickie, I. (2015). *Fixing Reference*, New York: Oxford University Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*, Cambridge, MA: MIT Press.
- Driver, J. and Baylis, G. C. (1989). Movement and visual attention: The spotlight metaphor breaks down, *Journal of Experimental Psychology: Human Perception and Performance* **15**: 448–456.
- Duncan, J. (1998). Converging levels of analysis in the cognitive neuroscience of visual attention, *Philosophical Transactions of the Royal Society* **353**(1373): 1307–1317.
- Duncan, J. and Humphreys, G. (1992). Beyond the search surface: Visual search and attentional engagement, *Journal of Experimental Psychology: Human Perception and Performance* **18**(2): 578–588.
- Egly, R., Driver, J. and Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects, *Journal of Experimental Psychology: General* **123**(2): 161–77.
- Elster, J. (2000). *Ulysses Unbound: Studies in Rationality, Precommitment, and Constraints*, Vol. 102, New York: Cambridge University Press, 121-138.
- Eriksen, B. A. and Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target in a non-search task, *Perception and Psychophysics* **16**: 143–149.

- Evans, C. O. (1970). *The Subject of Consciousness*, London: George Allen and Unwin.
- Evans, G. (1982). *The Varieties of Reference*, New York: Clarendon Press.
- Eysenck, M. W. and Derakshan, N. (2011). New perspectives in attentional control theory, *Personality and Individual Differences* **50**: 955–960.
- Eysenck, M. W., Derakshan, N., Santos, R. and Calvo, M. G. (2007). Anxiety and cognitive performance: Attentional control theory, *Emotion* **7**(2): 336–353.
- Fecteau, J. H. and Munoz, D. P. (2006). Saliency, relevance, and firing: A priority map for target selection, *Trends in Cognitive Sciences* **10**(8): 383–90.
- Fischer, J. M. (1994). *The Metaphysics of Free Will*, Malden MA: Blackwell Publishers.
- Fodor, J. (1983). *Modularity of Mind*, Cambridge, MA: MIT Press.
- Folk, C. L., Remington, R. W. and Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings, *Journal of Experimental Psychology: Human Perception and Performance* **18**(4): 1030–44.
- Fox, E., Russo, R., Bowles, R. and Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety?, *Journal of Experimental Psychology* **130**(4): 681–700.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person, *Journal of Philosophy* **68**(1): 5–20.
- Frankfurt, H. (1978). The problem of action, *American Philosophical Quarterly* **15**(2): 157–162.
- Frankfurt, H. (2002). Reply to michael bratman, *Contours of Agency: Essays on Themes from Harry Frankfurt*, Cambridge, MA: MIT Press.
- Frankfurt, H. G. (1976). Identification and externality, in A. O. Rorty (ed.), *The Identity of Persons*, Berkeley: University of California Press, pp. 239–252.
- Fridland, E. (2014). They've lost control: Reflections on skill, *Synthese* **191**(12): 2729–2750.
- Fridland, E. (2015). Knowing how: Problems and considerations, *European Journal of Philosophy* **23**(3): 703–727.
- Fridland, E. (2017a). Automatically minded, *Synthese* **194**(11): 4337–4363.

- Fridland, E. (2017b). Skill and motor control: Intelligence all the way down, *Philosophical Studies* **174**(6): 1539–1560.
- Frith, C. (2007). *Making Up the Mind: How the Brain Creates our Mental World*, Oxford: Blackwell.
- Frith, C. D., Blakemore, S. J. and Wolpert, D. M. (2000). Abnormalities in the awareness of action, *Philosophical Transactions of the Royal Society of London* **355**: 1771–99.
- Fuller, S. and Carrasco, M. (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue, *Vision Research* **46**(23): 4032–47.
- Ganson and Bronner (2013). Visual prominence and representationalism, *Philosophical Studies* **164**(2): 405–418.
- Gibson, B. S. and Kelsey, E. M. (1998). Stimulus-driven attentional capture is contingent on attentional set for displaywide visual features, *Journal of Experimental Psychology: Human Perception and Performance* **24**(3): 699–706.
- Gold, J. M., Fuller, R. L., Robinson, B. M., Braun, E. L. and Luck, S. J. (2007). Impaired top-down control of visual search in schizophrenia, *Schizophrenia Research* **94**(1-3): 148–55.
- Gottlieb, J. (2012). Attention, learning, and the value of information, *Neuron* **76**(18): 281–95.
- Henderson, J. M., Brockmole, J. R., Castelhana, M. S. and Mack, M. (2007). Visual saliency does not account for eye movements during visual search in real-world scenes, in R. P. G. van Gompel, M. H. Fischer, W. S. Murray and R. L. Hill (eds), *Eye Movements: A Window on Mind and Brain*, New York: Elsevier, pp. 539–62.
- Henderson, J. M., Malcolm, G. L. and Schandl, C. (2009). Searching in the dark: Cognitive relevance drives attention in real-world scenes, *Psychonomic Bulletin & Review* **16**(5): 850–856.
- Herwig, A. (2015). Linking perception and action by structure or process? toward an integrative perspective, *Neuroscience and Biobehavioral Reviews* **52**: 105–116.
- Horgan, T. (2015). Injecting the phenomenology of agency into the free will debate, in D. Shoemaker (ed.), *Oxford Studies in Agency and Responsibility*, Vol. 3, New York: Oxford University Press, pp. 34–61.

- Horgan, T. and Nichols, S. (2015). The zero point and i, in S. Miguens, G. Preyer and C. B. Morando (eds), *Prereflective consciousness: Sartre and contemporary philosophy of mind*, New York: Routledge, pp. 143–175.
- Horgan, T., Tienson, J. and Graham, G. (2003). The phenomenology of first-person agency, in S. Walter and H.-D. Heckmann (eds), *Physicalism and Mental Causation*, Charlottesville, VA: Imprint Academic, pp. 323–340.
- Hornsby, J. (2004). Agency and alienation, in M. D. Caro and D. Macarthur (eds), *Naturalism in Question*, Cambridge, MA: Harvard University Press., pp. 173–187.
- Husserl, E. (1983/1913). *Ideas pertaining to a pure phenomenology and to a phenomenological philosophy first book: General introduction to a pure phenomenology*, Leiden: Martinus Nijoff Publishers. Translated by F. Kersten.
- Israel, D., Perry, J. and Tutiya, S. (1993). Executions, motivations, and accomplishments, *The Philosophical Review* **102**(4): 515–540.
- Itti, L. and Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention, *Vision Research* **40**(10): 1489–506.
- Itti, L. and Koch, C. (2001). Computational modelling of visual attention, *Nature Reviews Neuroscience* **2**(3): 1–10.
- James, W. (1890/1950). *The Principles of Psychology*, Vol. 1, Dover.
- Jeannerod, M. (1997). *The Cognitive Neuroscience of Action*, Cambridge MA: Blackwell.
- Jeannerod, M. (2006). *Motor Cognition: What Actions Tell to the Self*, New York: Oxford University Press.
- Jennings, C. D. (2012). The subject of attention, *Synthese* **189**(3): 535–554.
- Jennings, C. D. (2015). Attention and perceptual organization, *Philosophical Studies* **172**: 1265–1278.
- Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement, in J. B. Long and A. D. Baddeley (eds), *Attention and Performance IX*, Hillsdale, NJ: Erlbaum, pp. 187–203.
- Jonides, J. and Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention, *Perception & Psychophysics* **43**: 346–54.

- Kastner, S. (2009). Attention, neural basis of, in P. W. T. Bayne, A. Cleeremans (ed.), *The Oxford Companion to Consciousness*, New York: Oxford University Press, pp. 72–78.
- Kelly, S. D. (2002). Merleau-pony on the body, *Ratio* **15**(4): 376–391.
- Kentridge, R. W. (2011). Attention without awareness: A brief review, in C. Mole, D. Smithes and W. Wu (eds), *Attention: Philosophical & psychological essays*, New York: Oxford University Press, pp. 228–246.
- Kentridge, R. W., Heywood, C. A. and Weiskrantz, L. (1997). Residual vision in multiple retinal locations: Implications for blindsight, *Journal of Cognitive Neuroscience* **9**: 191–202.
- Kimchi, R., Yeshurun, Y., Spehar, B. and Pirkner, Y. (2016). Perceptual organization, visual attention, and objecthood, *Vision Research* **126**: 34–51.
- Koch, C. and Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry, *Human Neurobiology* **4**: 219–27.
- Koralus, P. (2014). The erotetic theory of attention: Questions, focus, and distraction, *Mind & Language* **29**(1): 26–50.
- Kriegel, U. (2015). *The Varieties of Consciousness*, New York: Oxford University Press.
- Levy, N. (2017). Embodied savoir-faire: Knowledge-how requires motor representations, *Synthese* **194**(2): 511–530.
- Lewis, C. I. (1929). *Mind and the World Order: Outline of a Theory of Knowledge*, New York: Dover Publications.
- Liddle, P. M. (2006). Motivated attention and schizophrenia, in N. Sebanz and W. Prinz (eds), *Disorders of Volition*, Cambridge MA: MIT Press, pp. 193–205.
- Linnell, K. J., Humphreys, G. W., McIntyre, D. B., Laitinen, S. and Wing, A. M. (2005). Action modulates object-based selection, *Vision Research* pp. 2268–2286.
- Luthra, Y. (2016). Non-rational aspects of skilled agency, *Philosophical Studies* **173**(8): 2267–2289.
- Marcel, A. (2003). The sense of agency: Awareness and ownership of action, in J. Roessler and N. Eilan (eds), *Agency and Self-awareness: Issues in Philosophy and Psychology*, New York: Clarendon Press, pp. 48–93.
- Matthen, M. (2005). *Seeing, Doing, and Knowing: A Philosophical Theory of Sense Perception*, New York: Clarendon Press.

- Matthen, M. (2014). How to be sure: Sensory exploration and empirical certainty, *Philosophy and Phenomenological Research* **88**(1): 38–69.
- Mazzoni, P. and Krakauer, J. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation, *The Journal of Neuroscience* **26**(14): 3642–3645.
- McDougle, S., Ivry, R. and Taylor, J. (2016). Taking aim at the cognitive side of learning in sensorimotor adaptation tasks, *Trends in Cognitive Sciences* **20**(7): 535–544.
- Melden, A. I. (1961). *Free Action*, London: Routledge and Kegan Paul.
- Mele, A. (2003). *Motivation and Agency*, New York: Oxford University Press.
- Memelink, J. and Hommel, B. (2013). Intentional weighting: a basic principle in cognitive control, *Psychological Research* **77**(3): 249–259.
- Milner, D. A. and Goodale, M. A. (1995/2006). *The Visual Brain in Action*, Vol. 2nd, New York: Oxford University Press.
- Mole, C. (2009). Illusions, demonstratives, and the zombie action hypothesis, *Mind* **118**(472): 995–1011.
- Mole, C. (2011). *Attention is Cognitive Unison: An Essay in Philosophical Psychology*, New York: Oxford University Press.
- Mole, C. (2014). Attention to unseen objects, *Journal of Consciousness Studies* **21**(11-12): 41–56.
- Mole, C. (2015). Attention and cognitive penetration, in J. Zeimbekis and A. Raftopoulos (eds), *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, New York: Oxford University Press, pp. 218–238.
- Montero, B. G. (2016). *Thought in Action: Expertise and the Conscious Mind*, New York: Oxford University Press.
- Mylopoulos, M. (2012). Evaluating the case for the low-level approach to agentive awareness, *Philosophical Topics* **40**(2): 103–127.
- Mylopoulos, M. (2015). Agentive awareness is not sensory awareness, *Philosophical Studies* **3**(172): 761–80.
- Mylopoulos, M. (2017). A cognitive account of agentive awareness, *Mind & Language* **32**: 545–63.

- Nagel, T. (1974). What is it like to be a bat?, *The Philosophical Review* **83**(4): 435–450.
- Nagel, T. (1986). *The View from Nowhere*, New York: Oxford University Press.
- Nanay, B. (2010). Attention and perceptual content, *Analysis* **70**(2): 263–270.
- Neumann, O. (1987). Beyond capacity: A functional view of attention, in H. Heuer and A. F. Sanders (eds), *Perspectives on Perception and Action*, Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers., pp. 361–394.
- Pacherie, E. (2006). Toward a dynamic theory of intentions, in S. Pockett, W. P. Banks and S. Gallagher (eds), *Does Consciousness Cause Behaviour?*, Cambridge, MA: MIT Press, pp. 145–167.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework, *Cognition* **107**: 179–217.
- Pacherie, E. (2011). Nonconceptual representations for action and the limits of intentional control, *Social Psychology* **42**(1): 67–73.
- Papineau, D. (2013). In the zone, *Royal Institute of Philosophy Supplement* **73**: 175–196.
- Pautz, A. (2010). Why explain visual experience in terms of content?, in B. Nanay (ed.), *Perceiving the World*, New York: Oxford University Press, pp. 254–309.
- Peacocke, C. (1992). *A Study of Concepts*, Cambridge, MA: MIT Press.
- Peacocke, C. (2003). Action: Awareness, ownership, knowledge, in J. Roessler and N. Eilan (eds), *Agency and Self-awareness: Issues in Philosophy and Psychology*, New York: Clarendon Press, pp. 94–110.
- Peck, C. J., Jangraw, D. C., Suzuki, M., Efem, R. and Gottlieb, J. (2009). Reward modulates attention independently of action value in posterior parietal cortex, *The Journal of Neuroscience* **29**(36): 11182–91.
- Pereboom, D. (2014). The disappearing agent objection to event-causal libertarianism, *Philosophical Studies* **169**(1): 59–69.
- Perry, J. (1993). *The Problem of the Essential Indexical*, New York: Oxford University Press.
- Pessoa, L. (2013). *The Cognitive-Emotional Brain: From Interactions to Integration*, Cambridge, MA: MIT press.

- Posner, M. (1980). Orienting of attention, *Quarterly Journal of Experimental Psychology* **32**(1): 3–25.
- Posner, M. I. and Cohen, Y. (1984). Components of visual orienting, in H. Bouma and D. Bouwhuis (eds), *Attention and Performance X: Control of Language Processes*, Vol. 32, Hillsdale, NJ: Erlbaum, pp. 531–556.
- Posner, M. and Peterson, S. E. (1990). The attention system of the human brain, *Annual Review of Neuroscience* **13**: 25–42.
- Price, H. H. (1932). *Perception*, Methuen & Co.
- Prinz, J. (2004). *Gut reactions: A perceptual theory of emotion*, New York : Oxford University Press.
- Reid, T. (1872/2000). *The Works of Thomas Reid*, Elibron Classics.
- Riddoch, M. J. and Humphreys, G. W. . (2007). How to define an object: Evidence from the effects of action on perception and attention, *Mind and Language* **22**(5): 534–547.
- Robinson, T. E. and Berridge, K. C. (2008). The incentive sensitization theory of addiction: Some current issues, *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**: 3137–46.
- Roessler, J. (2000). Attention and the self: An appreciation of c.o. evans' "the subject of consciousness", *Journal of Consciousness Studies* **7**(5): 76–82.
- Roessler, J. (2003). Intentional action and self-awareness, in J. Roessler and N. Eilan (eds), *Agency and Self-awareness*, New York: Oxford University Press, pp. 382–405.
- Rothkopf, C. A., Ballard, D. H. and Hayhoe, M. M. (2007). Task and context determine where you look, *Journal of Vision* **7**(14): 1–20.
- Schlosser, M. E. (2011). Agency, ownership, and the standard theory, in J. H. Aguilar, A. A. Buckareff and K. Frankish (eds), *New waves in Philosophy of Action*, New York: Palgrave Macmillan, pp. 13–31.
- Schneider, K. A. and Komlos, M. (2011). Attention alters decision criteria but not appearance: a reanalysis of anton-erxleben, abrams, and carrasco (2010), *Journal of Vision* **11**(13): 1–10.
- Schneider, W. and Deubel, H. (2002). Selection for perception and selection for spatial motor action are coupled by visual attention: a review of recent findings and new evidence

- from stimulus-driven saccade control, in W. Prinz and B. Hommel (eds), *Attention and Performance XIX: Common Mechanisms in Perception and Action*, New York: Oxford University Press, pp. 609–627.
- Schwenkler, J. (2014). Vision, self-location, and the phenomenology of 'the point of view', *Nous* **48**(1): 137–155.
- Searle, J. (1983). *Intentionality: An Essay in the Philosophy of Mind*, Cambridge, MA: Cambridge University Press.
- Serences, J. T. and Yantis, S. (2006). Selective visual attention and perceptual coherence, *Trends in Cognitive Sciences* **10**(1): 38–45.
- Shepherd, J. (2016a). Conscious action/zombie action, *Nous* **50**(2): 419–444.
- Shepherd, J. (2016b). Conscious control over action, *Mind & Language* **30**(3): 320–44.
- Shepherd, J. (2017a). The experience of acting and the structure of experience, *Journal of Philosophy* **114**(8): 422–448.
- Shepherd, J. (2017b). Halfhearted action and control, *Ergo* **14**(9): 259–76.
- Shmuelof, L., Krakauer, J. and Mazzoni, P. (2012). How is a motor skill learned? change and invariance at the levels of task success and trajectory control, *Journal of Neurophysiology* **108**: 578–594.
- Shulman, G. L. and Corbetta, M. (2012). Two attentional networks: Identification and function within a larger cognitive architecture, in M. I. Posner and L. Huang (eds), *Cognitive Neuroscience of Attention, Second Edition*, New York: Guilford Press, p. 113–128.
- Siegel, S. (2006). Subject and object in the contents of visual experience, *The Philosophical Review* **115**(3): 355–388.
- Smith, D. E. and Schenk, T. (2012). The premotor theory of attention: Time to move on?, *Neuropsychologia* **50**(6): 1104–1114.
- Smithies, D. (2011). Attention is rational-access consciousness, in C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological Essays*, New York: Oxford University Press, pp. 247–73.
- Speaks, J. (2010). Attention and intentionalism, *Philosophical Quarterly* **60**(239): 325–42.

- Spence, C. J. and Driver, J. (1994). Covert spatial orienting in audition: Exogenous and endogenous mechanisms, *Journal of Experimental Psychology: Human Perception and Performance* **20**(3): 555–574.
- Spence, C. and Santangelo, V. (2009). Capturing spatial attention with multisensory cues: A review, *Hearing Research* **258**(1): 134–42.
- Spence, S. A., Brooks, D. J., Hirsch, S. R., Liddle, P. F., Meehan, J. and Grasby, P. M. (1997). A pet study of voluntary movement in schizophrenic patients experiencing passivity phenomena (delusions of alien control), *Brain* **120**(11): 1997–2011.
- Stanley, J. (2011). *Know how*, New York Oxford University Press.
- Stanley, J. and Krakauer, J. W. (2013). Motor skill depends on knowledge of facts, *Frontiers in Human Neuroscience* **7**(503): 1–11.
- Stanley, J. and Williamson, T. (2001). Knowing how, *Journal of Philosophy* **98**: 411–444.
- Stazicker, J. (2011a). *Attention and the Indeterminacy of Visual Experience*, PhD thesis, University of California, Berkeley.
- Stazicker, J. (2011b). Attention, visual consciousness, and indeterminacy, *Mind & Language* **26**(2): 156–184.
- Steward, H. (2012). *A Metaphysics for Freedom*, New York: Oxford University Press.
- Stinson, C. (2009). Searching for the source of executive attention, *Psyche* **15**(1): 137–154.
- Styles, E. A. (2006). *The Psychology of Attention*, 2nd edn, Psychology Press: Taylor and Francis Group.
- Tatler, B. W., Hayhoe, M. M., Land, M. F. and Ballard, D. H. (2011). Eye guidance in natural vision: Reinterpreting salience, *Journal of Vision* **11**(5): 1–23.
- Tipper, S. P., Howard, L. A. and Houghton, G. (1998). Action-based mechanisms of attention, *Philosophical Transactions of the Royal Society of London* **353**: 1385–1393.
- Tipper, S. P., Lortie, C. and Baylis, G. C. (1992). Selective reaching: Evidence for action-centered attention, *Journal of Experimental Psychology: Human Perception and Performance* **18**(4): 891–905.
- Todd, R. M., Cunningham, W. A., Anderson, A. K. and Thompson, E. (2012). Affect-biased attention as emotion regulation, *Trends in Cognitive Sciences* **16**(7): 365–72.

- Todd, R. and Manigold, G. (2018). Implicit guidance of attention: The priority state space framework, *Cortex* **102**: 121–138.
- Todorov, E. (2004). Optimality principles in sensorimotor control, *Nature Neuroscience Review* **7**(9): 907–15.
- Todorov, E. and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination, *Nature Neuroscience* **5**(11): 1226–1235.
- Treisman, A. M. (1964). The effect of irrelevant material on the efficiency of selective listening, *American Journal of Psychology* **77**: 533–46.
- Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention, *Cognitive Psychology* **12**: 97–136.
- Velleman, J. D. (1992). What happens when someone acts?, *Mind* **101**(403): 461–481.
- Velleman, J. D. (2000a). The aim of belief, *The Possibility of Practical Reflection*, New York: Oxford University Press, pp. 244–81.
- Velleman, J. D. (2000b). *The Possibility of Practical Reason*, New York: Clarendon Press.
- Velleman, J. D. (2008). The way of the wanton, in K. Atkins and C. MacKenzie (eds), *Practical Identity and Narrative Agency*, New York: Routledge, pp. 169–192.
- Vine, S. J., Lee, D., Moore, L. J. and Wilson, M. R. (2013). Quiet eye and choking: Online control breaks down at the point of performance failure, *Medicine and Science in Sports and Exercise* **45**(10): 1988–94.
- Vuilleumier, P. (2005). How brains beware: Neural mechanisms of emotional attention, *Trends in Cognitive Sciences* **9**(12): 585–94.
- Wagemans, J., Elder, J. H., Palmer, S. E., Peterson, M. A. and Singh, M. (2012). A century of gestalt psychology in visual perception: I. perceptual grouping and figure-ground organization, *Psychological Bulletin* **138**(6): 1172–1217.
- Watzl, S. (2010). *The Significance of Attention*, PhD thesis, University of Columbia.
- Watzl, S. (2011). Attention as structuring of the stream of consciousness, in C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological Essays*, New York: Oxford University Press, pp. 145–173.
- Watzl, S. (2017). *Structuring Mind: The Nature of Attention and How it Shapes Consciousness*, Oxford University Press.

- Watzl, S. (forthcoming). Can intentionalism explain how attention affects appearances?, in A. Pautz and D. Stoljar (eds), *Themes from Ned Block*, Cambridge, MA: MIT Press.
- Watzl, S. (in preparation). Attention and the subjective point of view.
- White, A. (1964). *Attention*, Oxford: Basil Blackwell.
- Williams, B. (1973). Deciding to believe, in B. Williams (ed.), *Problems of the Self*, Cambridge, MA: Cambridge University Press, pp. 136–51.
- Wu, W. (2011a). Attention as selection for action, in C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological essays*, New York: Oxford University Press, pp. 97–116.
- Wu, W. (2011b). Confronting many-many problems: Attention and agentic control, *Nous* 45(1): 50–76.
- Wu, W. (2011c). What is conscious attention?, *Philosophy and Phenomenological Research* 82(1): 93–120.
- Wu, W. (2014a). Against division: Consciousness, information and the visual streams, *Mind & Language* 29(4): 383–406.
- Wu, W. (2014b). *Attention*, New York: Routledge.
- Wu, W. (2016). Experts and deviants: The story of agential control, *Philosophy and Phenomenological Research* 93(1): 101–126.
- Wu, W. (2017). Shaking up the mind’s ground floor: The cognitive penetration of visual attention, *Journal of Philosophy* 114(1): 5–32.
- Yantis, S. and Jonides, J. (1984). Abrupt visual onsets and selective attention: Evidence from visual search, *Journal of Experimental Psychology: Human Perception and Performance* 10: 601–21.
- Yeshurun, Y. and Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution, *Nature* 396: 72–5.