A Question-Sensitive Theory of Intention

Bob Beddor & Simon Goldstein

Abstract

This paper develops a question-sensitive theory of intention. We show that this theory explains some puzzling closure properties of intention. In particular, it can be used to explain why one is rationally required to intend the means to one's ends, even though one is not rationally required to intend all the foreseen consequences of one's intended actions. It also explains why rational intention is not always closed under logical implication, and why one can only intend outcomes that one believes to be under one's control.

1 Introduction

Intending one thing often commits you to intending another. Suppose you intend to brew some coffee and sip it while reading this paper. This rationally commits you to intend many of the things you take to be material consequences of this action (for example, you are rationally required to intend to take whatever steps you deem necessary for coffee-brewing). But you are not committed to intending all such believed consequences (for example, you need not intend for this paper to have been written, or for your heart-rate to become elevated from the coffee). This paper seeks to develop a systematic theory of rational intention that explains the conditions under which one intention commits you to another.

Our theory draws on a rich body of recent research that takes certain propositional attitudes to be question-sensitive. Yalcin 2011, 2018, Drucker 2020 and Hoek forthcoming argue that agents do not believe propositions *simpliciter*; rather, they only believe propositions relative to a question that is under consideration. This sort of view offers to explain various closure failures for belief, thereby making progress on the problem of logical omniscience.[1] In a similar vein,

[1] For development of a question-sensitive theory of mathematical belief, see Pérez Carballo 2016. See also Berto 2018, 2019 for question-sensitive treatments of imagination and belief revision, Holguín 2022 for a question-sensitive treatment of belief which permits belief in low-probability events, and Mandelkern et al. 2017 for a question-sensitive treatment of agentive modals. A complementary thread of research explores the idea that conversations unfold against the backdrop of a question under discussion. See Roberts 1996/2012 for an overview.

Schaffer 2005, 2007, Yablo 2014, and Hawke 2016 suggest that a question-sensitive theory of knowledge explains a range of linguistic data and makes headway on long-standing epistemological puzzles, for example, the skeptical paradox. And Phillips-Brown 2018 extends these ideas to defend a question-sensitive theory of desire.

This paper contributes to this research program by developing a question-sensitive theory of rational intention. According to our proposal, intentions are formed against the backdrop of a practical question: *What to do?* Rational constraints on intention derive—at least in part—from rational constraints on how to address practical questions. We show that this basic idea can be developed in a way that explains some puzzling closure properties of intention. In particular, it explains why one is rationally required to intend the means to one's ends, even though one is not rationally required to intend all the foreseen consequences of one's intended actions. It also explains why rational intention is not always closed under logical implication, and why one can only intend outcomes that one believes to be under one's control.

Here's the game plan. §2 introduces some distinctive puzzles involving closure for intention. §3 reviews an off-the-shelf semantics for intention reports in the tradition of Hintikka 1962, and explains why it fails to resolves our puzzles. This failure motivates the development of our positive account, which is developed in §§4-6. Finally, §7 uses our theory to shed light on the relation between intention and desire.

2 Two Puzzles of Closure

In epistemology, much ink has been spilled over the extent to which knowledge obeys closure constraints. On the one hand, it is a platitude that deduction is often a way of extending our knowledge. On the other hand, unrestricted closure principles face counterexamples.[2] This creates a two-fold challenge for epistemologists: First, to articulate a restricted closure principle that describes when we can know the consequences of what we know. Second, to develop theories of knowledge that validate this restricted closure principle.

By contrast, closure puzzles involving intention have attracted much less attention. Here, we will introduce two such puzzles. The first concerns the extent to which rational intention is closed under logical consequence. The

[2] Perhaps the most famous arise from the skeptical paradox and the lottery puzzle. See Dretske 1970; Nozick 1981; Hawthorne 2004, among many others.

second concerns the extent to which rational intention is closed under believed consequence.

Start with logical consequence. On the one hand, we are often rationally required to intend the obvious logical consequences of our intentions. For example, the following seem to ascribe patently irrational states of mind:

(1) ?? Ana intends to ski at Zermatt and Chamonix this winter, but she doesn't intend to ski at Chamonix this winter.

(2) ?? Ana intends to ski at Zermatt and Chamonix this winter, but she doesn't intend ski at either Zermatt or Chamonix this winter.

These data generate some pressure to embrace a single-premise logical closure requirement on intention. Where $\models$ represents logical consequence:

Unrestricted Logical Closure If a rational agent A intends p, and $p \models q$, then A also intend q.[3]

On the other hand, Unrestricted Logical Closure faces at least two types of counterexample. The first—familiar from discussion of closure principles on knowledge and rational belief—comes from cases where an agent lacks the conceptual resources to grasp the relevant entailment. Consider William III, who lacked the concept of nuclear technology. (3-a) could be true, even though (3-b) is false:

(3) a. William III intended to avoid war with France. $\not\models$

[3] Here we assume that the objects of intention are propositions. While this assumption is fairly common, it is not uncontroversial. After all, the most natural way of ascribing intentions—as exemplified in (1) and (2)—is via the form, *A intends to VP*, which might be taken to show that the objects of intention are actions or properties. However, there is a natural response to this worry (Davis 1984; Ferrero 2013). When we say *Ana intends to ski this winter*, this is equivalent to an ascription of the form, *Ana$_i$ intends that she$_i$ skis this winter*. In favor of this approach, note that while ascriptions of intentions with that-clauses are somewhat uncommon, they are certainly coherent, and can easily be found in the wild. Some examples from the Corpus of Contemporary American English (Davies 2008-):

> 'He intends that there will only be two classes of people in this world - the elite nobility. . . and the serfs. . . '

> 'Joe Kearney did not live a futile life; I intend that he shall not have died a futile death.'

By taking the objects of intention to be propositions, we can capture the full range of intention ascriptions, both those of the form, *A intends that p* and those of the form, *A intends to VP*.

For those who find this reply unconvincing, a more concessive response is also possible. Even if we grant that the objects of intention are actions, it's easy enough to convert talk of intending a proposition into talk of intending an action. Whenever we say, *A intends p*, simply replace this with, *A intends to bring about p*.

b. William III intended to avoid nuclear war with France.[4]

A second type of counterexample is not shared by knowledge or rational belief. We cannot rationally intend outcomes that we believe to be outside of our control, even if these outcomes are obvious logical consequences of something we intend. For example:

(4) a. The police intend to arrest whoever committed the murder. $\not\models$
b. The police intend for there to have been a murder.

(5) a. John intends to die an honorable death. $\not\models$
b. John intends to be mortal.

Here the failure of closure is not due to any cognitive or conceptual limitations.[5]

This gives rise to what we'll call the 'The Puzzle of Logical Consequence': the puzzle of developing a theory of intention that explains why we are often required to intend the logical consequences of our intentions, while also accounting for the two classes of counterexamples to Unrestricted Logical Closure.

A second puzzle concerns the extent to which intention is closed under *believed* consequence. On the one hand, it is widely thought that we are under rational pressure to intend the believed necessary means to our ends:

MEANS-ENDS COHERENCE If an agent A is certain that q is one of the necessary means of bringing about p, then it is irrational for A to intend p without also intending q.[6]

To illustrate the plausibility of this requirement, suppose Barry is certain that in order to make a delicious feast, he must first stop by the grocery store. Then the following ascribes a defective planning state to Barry:

(6) ?? Barry intends to cook a delicious feast, but he doesn't intend to stop by the grocery store.

At the same time, common sense holds that we are not required to intend every foreseen consequence of our intended ends:

[4] Cf. Stalnaker 1984 and Yalcin 2018 for discussion of an analogous closure failure for belief.

[5] We discuss whether similar closure failures arise for desire in §7.

[6] See e.g., Bratman 1987; Broome 1999; Kolodny and Brunero 2020. Throughout we will assume that the relevant means are not just any necessary conditions for the fulfillment of the intention; rather, they are actions that the agent believes to be under their control. (Otherwise Means-Ends Coherence would implausibly predict that someone who intends to run a marathon also intends for oxygenated blood to enter their heart, which would conflict with the observation that we can only intend outcomes that we believe to be under our control.)

Foreseeing Without Intending A rational agent A can intend p, and be certain that p will lead to q, while failing to intend to q.

In the philosophy of action and ethics literature, this principle is often motivated by cases like the Tactical Bomber, who bombs an enemy munitions factory, intending to decrease the enemy's supply of munitions. In doing so, they foresee that the bombing will kill all nearby civilians, though they regret that their action will have this consequence. A common intuition is that:

(7) The Tactical Bomber intends to decrease the enemy's munitions, but they do not intend to kill all the nearby civilians.

This intuition has shaped a long tradition in ethics. For example, according to the doctrine of double effect, there is an important moral difference between Tactical Bomber's action and that of a terrorist bomber, who drops the same bomb with the goal of killing all nearby civilians.[7,8]

Taken together, Means-Ends Coherence and Foreseeing Without Intending reveal an important asymmetry between the means and the effects of our intended ends. To illustrate with another example, consider Carol the carouser. She might rationally intend to get drunk without intending to get a hangover, even if she foresees that getting drunk will lead to a hangover. But she cannot rationally intend to get drunk without intending to buy alcohol, provided she believes that buying alcohol is a necessary means to getting drunk.

This asymmetry gives rise to what we'll call, 'The Puzzle of Believed Consequence': the problem of developing a theory of intention that explains why we are rationally required to intend the believed means to our ends, even though we are not required to intend the believed consequences of our ends.

These two puzzles are distinct. Still, we might well hope for a unified solution to these two puzzles. And this is what we aim to deliver.

[7] For discussion, see, among others, Bennett 1980; Quinn 1989; McMahan 2009; Nelkin and Rickless 2014; Tadros 2015. See McIntyre 2019 for a general overview.

[8] As a referee helpfully points out, the Tactical Bomber case bears a resemblance to the 'Knobe Effect': the observation that moral valence of a side-effect influences our judgments about whether someone brought about that side-effect intentionally (Knobe 2003). An interesting question is whether the the account of Foreseeing Without Intending developed here could be extended to shed light on the Knobe Effect.

3 Hintikka Semantics

A natural starting point for theorizing about propositional attitudes is a Hintikka semantics (Hintikka 1962), where attitude verbs are analyzed as modal operators. For example, A *believes* p is true iff p holds in all of A's doxastic alternatives—that is, all the worlds consistent with what A believes. Similarly, A *knows* p is true iff p holds in all of A's epistemic alternatives—that is, all of the worlds consistent with what A knows.

Extended to intention, this analysis says that A *intends* p is true iff p holds in all of A's conative alternatives—that is, all of the worlds consistent with what A intends.[9] Where $Con(w)$ represents the conative alternatives of agent A at world w:

$$[\![A \text{ intends } p]\!]^w = 1 \text{ iff } Con(w) \subseteq [\![p]\!]$$

This semantics does not provide a reductive account of intention, any more than a Hintikka semantics for *believes* gives a reductive account of belief. Those with reductive ambitions could supplement it with a reductive analysis of the conative alternatives, to be cashed out in dispositionalist or functionalist terms. (In §5 we explore a few candidate reductions.) Still, even without such an account, a Hintikka semantics generates a number of substantive predictions about the logic of intention. We now investigate these predictions regarding our two closure puzzles.

Start with the Puzzle of Logical Consequence. A familiar feature of a Hintikka semantics is that it validates Unrestricted Logical Closure. After all, if p holds throughout your conative alternatives, and p entails q, then q also holds throughout your conative alternatives. Thus a Hintikka semantics explains why it often sounds odd to assert that an agent does not intend the obvious consequences of their intended ends, as revealed by (1) and (2). But it does so at the cost of over-predicting instances of closure. In particular, it does not explain our two classes of counterexamples (e.g., (3-a) $\not\models$ (3-b); (4-a) $\not\models$ (4-b); (5-a) $\not\models$(5-b)).

Turn next to the Puzzle of Believed Consequence. As it stands, a Hintikka semantics does not make any predictions concerning whether we are rationally required to intend either the believed means to, or the believed consequences of, our ends. However, we could impose various constraints on conative alternatives

[9] More precisely: A *intends* p is true at time t iff p holds in all of A's conative alternatives at t. For the sake of notational simplicity, we will omit reference to times throughout.

in order to generate such predictions. For example, in order to validate Means-Ends Coherence, we could impose the constraint that a rational agent's conative alternatives must be a subset of their doxastic alternatives. Where $Dox(w)$ represents the agent's doxastic alternatives:

Belief Containment $Con(w) \subseteq Dox(w)$[10]

Belief Containment suffices to validate Means-Ends Coherence. To see this, recall Carol the carouser, who intends to get drunk. Since she intends to get drunk, she does so at every conative alternative. By Belief Containment, her conative alternatives are a subset of her doxastic alternatives. Now suppose that Carol believes that buying alcohol is necessary for getting drunk. In that case, Carol buys alcohol at every doxastic alternative where she gets drunk, and hence at every conative alternative, from which it follows that she intends to buy alcohol.

Unfortunately, while Belief Containment validates Means-Ends Coherence, it rules out Foreseeing without Intending. Carol not only believes that buying alcohol is necessary for getting drunk. She also believes that getting drunk ensures getting a hangover. So at every doxastic alternative where she gets drunk, a hangover awaits. By Belief Containment, Carol gets a hangover at every conative alternative. Consequently, she intends to get a hangover.

At this point, some might observe that Belief Containment is independently implausible, since it predicts that whatever you believe you also intend. But this is absurd. One obvious counterexample comes from the control requirement on intention. You might believe something that you don't intend, since you might recognize that you have no control over the matter, e.g.:

(8) a. John believes that the Western Roman Empire fell in 476. $\not\models$
b. John intends that the Western Roman Empire fell in 476.

In light of this observation, some might suggest that we should replace the Belief Containment with the weaker principle that your intentions are merely consistent with what you believe:

Belief Overlap $Con(w) \cap Dox(w) \neq \emptyset$.[11]

[10] Harris 2022 defends a Hintikka semantics for intention that imposes this constraint. See von Fintel 1999 for a development of Belief Containment in the case of desire reports, along with strategies for addressing the implausible consequences of this constraint.

[11] Building on Condoravdi and Lauer 2016, Grano 2017 develops a semantics for intention

Belief Overlap allows for the possibility of foreseeing that something will occur without intending it, since it allows that the conative alternatives are not contained within the doxastic alternatives. But for the purposes of our puzzle, Belief Overlap is too weak, since it does not ensure that Means-Ends Coherence holds.

So when it comes to the Puzzle of Believed Consequence, the Hintikka semantics faces a dilemma. On the one hand, it can be used to validate Means-Ends Coherence, provided we supplement it with Belief Containment. But this would invalidate Foreseeing Without Intending. On the other hand, we could avoid this consequence by weakening Belief Containment to Belief Overlap, but then we no longer validate Means-Ends Coherence.

At this point, some might suggest abandoning the Hintikka framework altogether. One natural place to cast about for an alternative is in the literature on desire attributions, which contains a range of proposals tailored to make room for closure failures. We think this reaction is too hasty. In §§4-6, we lay out our preferred approach, which keeps the basic apparatus of a Hintikka semantics and Belief Containment, but uses question-sensitivity to address the problems noted above. Having developed our account, §7 contrasts intention with desire, and shows that extant theories of desire do not generalize to explain the closure profile of intention.

4 Intention is Question-Sensitive

When someone is forming an intention, they are settling on a course of action. Or, to put it another way, they are settling on an answer to the practical question: *What to do?*

This much seems like a platitude. We propose to integrate this platitude into a semantics for intention ascriptions. On our analysis, intention ascriptions are sensitive to a practical question. Rational constraints on intention derive partly from rational constraints on how to address the practical question at issue.

that validates Belief Overlap. Intentions are understood in terms of an agent's 'effective' preference orderings over propositions. These effective preferences are action-guiding, and are required to be consistent with one's beliefs. Grano 2017 also requires that for any rational agent, the set of maximally ranked effective preferences are collectively consistent with their beliefs. The result is that for rational agents, Grano endorses a type of Hintikka semantics for intention, which universally quantifies over the worlds consistent with the agent's highest ranked effective preferences. (A wrinkle: Grano 2017 also requires that an agent intends p only if the highest ranked worlds imply a distinct proposition from p: namely, that the agent A is responsible for p. For present purposes, we will set aside this complication.)

4.1 The Semantics

Let's start by formulating the notion of a 'practical question' more precisely. When an agent is deliberating about what to do, they are using practical reasoning to decide between alternative courses of action. This suggests that we can model the practical question, *What to do?* as a partition of some region of logical space into a set of propositions describing various actions that the agent is deliberating between at a particular time.[12,13]

Which region, exactly? We presumably want to exclude worlds where the agent does not exist, or occupies completely different circumstances. One option would be to let the region be the set of worlds consistent with the agent's beliefs, or perhaps the set of worlds consistent with some privileged subset of the agent's beliefs—e.g. their beliefs about the outcomes they can affect. Alternatively, we could let the region be some set of worlds that resembles the agent's world in certain respects, akin to a circumstantial modal base (cf. Kratzer 1981).

This way of representing practical questions has close ties to decision theory, where it is common to model an agent's decision using a set of mutually exclusive actions available to the agent. Given the familiarity of this decision theoretic apparatus, we will refer to such a partition as the agent's 'decision problem.'[14] In §5.1 we consider some choice points for fleshing out the notion of a decision problem in more detail. But for present purposes, it will suffice to illustrate the basic idea with an example.

Suppose Barry is deciding what to cook for dinner. Then as a first pass, Barry's decision problem is a partition of his doxastic alternatives into all of the different culinary courses he is deciding between:

$$\{\textsf{cook hamburgers}, \textsf{cook halibut}, \textsf{cook risotto}, \ldots \textsf{cook ortolan}, \textsf{cook escargot}\}$$

Equipped with the notion of a decision problem, we now say what it is for a proposition to be defined on a decision problem. Say that a proposition is defined on a question Q, given some background restriction of logical space R, just in case the conjunction of the proposition and R is logically equivalent to

[12] This is in line with a rich tradition of research on questions (Hamblin 1958; Groenendijk and Stokhof 1984; Lewis 1988a,b). We explore the interactions between these models of questions and our semantics in §6.

[13] Here we leave open whether to understand deliberation—or planning more broadly—as a conscious or dispositional state. See Bratman 1987 for relevant discussion.

[14] See Cariani 2013 for the view that the semantic values of deontic modals are relativized to the options available to the agent. We should note that whereas decision problems are sometimes taken to include an agent's utilities and a representation of states of the world, we use the label more restrictively to refer to just the set of available actions.

some union of cells from Q, restricted to R. As a special case where Q is the agent's decision problem and R is the relevant restriction, we get:

DEFINEDNESS p is defined on decision problem D iff $\exists D' \subseteq D$ where $p \cap R = (\bigcup D') \cap R$.

To illustrate, consider Figure 1.

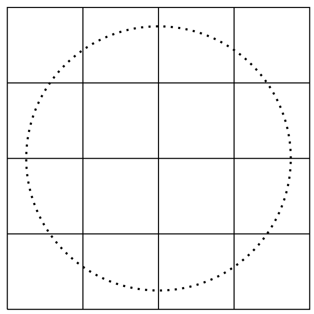

Figure 1: Definedness

Here the agent's decision problem partitions some relevant set of worlds into 16 cells, representing courses of action that the agent is deliberating between. A proposition is defined when it is some union of these cells. So the dotted region represents an undefined proposition relative to this decision problem.

To make this more concrete, return to Barry's culinary quandary. We can think of the 16 cells as the 16 different culinary options Barry is deciding between. The proposition, *Barry cooks hamburgers* is defined on his decision problem, as is the proposition, *Barry cooks hamburgers or Barry cooks risotto or Barry cooks ortolan*. After all, both of these propositions are equivalent to some union of the different options available to him.

We propose enriching the Hintikka semantics from §3 with the constraint that an agent intends p only if p is defined on their decision problem. As before, let $Con(w)$ represent A's conative alternatives at w. And let D represent A's decision problem at w. We propose:

$$[\![A \text{ intends } p]\!]^w = 1 \text{ iff (i) } Con(w) \subseteq [\![p]\!] \text{ and } (ii)\ [\![p]\!] \text{ is defined on } D.$$

A natural question for this approach concerns how to handle simultaneous intentions. In any realistic scenario, Barry will have many intentions at the same time. For example, Barry might not just have intentions regarding tonight's dinner, but also have intentions regarding tomorrow's lunch.

One option would be to associate every intention with a distinct decision problem. However, this approach faces an important drawback. Rational agents are under pressure to integrate their intentions in order to come up with a coherent course of action. Barry's dinner intentions may constrain the morrow's luncheoning options. For this reason, a better option is to use a single, more fine-grained decision problem that reflects all of the options that an agent is deliberating between on a given occasion. On this view, Barry's decision problem is better understood as a more fine-grained set, e.g.:

$$\{\textsf{burgers for dinner \& halibut for lunch}, \textsf{halibut for dinner \& burgers for lunch}, \ldots\}$$

If Barry has further intentions at the same time—say, concerning which movie to watch after dinner—then his decision problem will be more fine-grained still.

A principled reason for this way of handling simultaneous intentions comes from the logic of intention. As we will see in §6, an important constraint on the logic of intentions is that the intentions of rational agents agglomerate. If we associated each intention with a separate decision problem, there would be no guarantee that intentions agglomerate. But by using a single decision problem that integrates all of an agent's practical alternatives, we can easily validate this requirement.

Having laid out our semantics, let us now apply it to our two closure puzzles.[15]

4.2 Solving the Puzzles

Solving the Puzzle of Logical Consequence

Our definedness condition invalidates Unrestricted Logical Closure. To see this, consider Figure 2. Here, the agent's decision problem is again represented by a grid of 16 squares. Their conative alternatives $Con(w)$ are represented by the innermost circle. Note that the agent's conative alternatives are not

[15] There are parallels between our theory and other proposals in the literature. For example, Yalcin 2011, 2018, Berto 2019, and Hoek forthcoming all model belief in terms of a function from questions to sets of worlds. In addition to the obvious difference that these theories are focused on belief rather than intention, it is worth noting that these theories impose relatively few constraints on which questions are used to determine an agent's question-relative belief state. By contrast, one of our central aims is to specify substantive and precise constraints on admissible practical questions in order to explain the distinctive closure properties of intention. Another point of comparison is Phillips-Brown 2018, who uses questions to impose a definedness constraint on desire attributions.

intended, because they are not defined on their decision problem. The bold square represents the strongest action, p, that is defined on their decision problem and also intended. p^- is logically weaker than p, because it is true at strictly more worlds. But unlike p, p^- is not defined relative to the agent's decision problem. So they intend p without intending p^-, even though p implies p^-.

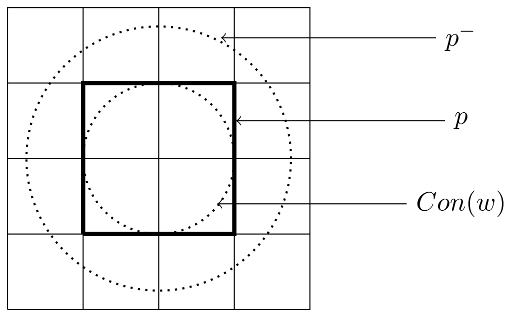

Figure 2: Closure failures

At the same time, our theory validates logical closure whenever the conclusion is defined on the agent's decision problem:

Restricted Logical Closure If (i) a rational agent A intends p, (ii) $p \models q$, and (iii) q is defined on A's decision problem, then A intends q.

This suggests that the current semantics has the potential to solve the Puzzle of Logical Consequence. On the one hand, our semantics predicts that we often intend the logical consequences of our intentions; we do so whenever these consequences are also defined on our decision problem. On the other hand, it allows that we don't always intend the logical consequences of our intentions, since sometimes these consequences are undefined for us.

Having seen the general shape of the solution, let's consider the counterexamples to Unrestricted Logical Closure in more detail. The first counterexample involved agents who lack the concepts required to grasp the logical consequence of their intentions:

(3) a. William III intended to avoid war with France. $\not\models$
b. William III intended to avoid nuclear war with France.

This invalidity falls immediately out of our semantics, given our understanding of a decision problem. A decision problem represents the actions that an agent is deliberating about whether to perform. But you cannot deliberate about whether to perform an action if you lack the conceptual tools to cognitively represent that action. So avoiding nuclear war does not belong to William III's decision problem. Figure 2 can model this case, with p representing the worlds where William III avoids war, and p^- representing the worlds where he avoids nuclear war.

The second counterexample to Unrestricted Logical Closure involved agents who believe the logical consequence of their action to be outside of their control:

(4) a. The police intend to arrest whoever committed the murder. $\not\models$
b. The police intend there to have been a murder.

(5) a. John intends to die an honorable death. $\not\models$
b. John intends to be mortal.

The invalidity of these inferences is also an immediate consequence of our understanding of a decision problem. You cannot rationally deliberate about whether to perform an action if you are convinced that this action is out of your control. That is:

Control Constraint $\forall p : p \in D$ only if the agent believes that p might be under their control.

So our theory avoids both classes of counterexample to the Unrestricted Logical Closure.

Solving the Puzzle of Believed Consequence

Turn next to the Puzzle of Believed Consequence. Any solution to this puzzle will validate Means-End Coherence. Without further elaboration, our theory does not do this. In order to remedy this, we propose enriching our theory with two further constraints. First, we adopt Belief Containment, requiring that $Con(w)$ is included in $Dox(w)$. In §3, we saw that in the presence of a simple Hintikka semantics, Belief Containment has the absurd implication that an agent intends whatever she believes. This is not a consequence of our framework. To see why, consider Figure 3.

Here we enrich our previous example with a representation of the agent's belief state $Dox(w)$, the dotted set of worlds nested between $Con(w)$ and p^-. In this example, any claim that the agent believes is guaranteed to hold throughout $Con(w)$. p^- is one such claim. But not every claim believed by the agent is

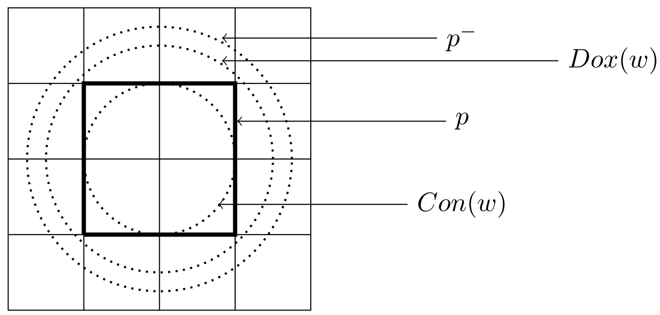

Figure 3: Foreseeing without intending

intended, since not every such claim is defined on the agent's decision problem. Our theory thus explains the invalidity of inferences like (8):

(8) a. John believes that the Western Roman Empire fell in 476. $\not\models$
b. John intends that the Western Roman Empire fell in 476.

After all, the Control Constraint ensures that *The Western Roman Empire fell in 476* is not defined on John's decision problem.

Next, we assume that decision problems are sensitive to causal structure. In particular, we require that whenever the agent believes that q is a necessary means for p and p is defined on the agent's decision problem, then q is also defined on that problem. That is:

Causal Constraint If both:

1. p is defined on D,
2. the agent believes q is a necessary means to p,

then q is also defined on D.

This constraint is not purely stipulative; it falls out from our understanding of decision problems. Recall that the cells in a decision problem represent actions that the agent is deliberating over. So the Causal Constraint really amounts to the plausible idea that if you are seriously considering performing some action p, and you believe that in order to perform that action, you must first perform some other action q which is also under your control, then you are under rational pressure to seriously consider performing q.

By combining these two ingredients—Belief Containment and the Causal Constraint—our theory validates Means-Ends Coherence. Suppose an agent intends p, and that they also believe that q is a necessary means to p. By our theory, it follows that both (i) p holds throughout their conative alternatives, (ii) p is defined on their decision problem. From (i) and Belief Containment, it follows that q holds throughout their conative alternatives. From (ii) and the Causal Constraint, it follows that q is also defined on their decision problem. So the agent intends q.

Crucially, while we impose the Causal Constraint, we do not impose the further requirement that the foreseen consequences of p are also defined on the decision problem. Thus we also validate Foreseeing Without Intending.

To illustrate, recall Carol the carouser. She intends to drink, and she also believes that buying alcohol is a necessary means to achieving this end. Now consider two possible ways of representing Carol's decision problem. The first is a decision problem that respects the Causal Constraint:

$$\{\textsf{buy \& drink}, \textsf{buy \& no drink}, \textsf{no buy \& drink}, \textsf{no buy \& no drink}\}$$

The propositions, *Carol drinks*, *Carol buys*, and *Carol buys and drinks* are all defined on this decision problem. However, the proposition *Carol gets a hangover* is not defined on this decision problem, since it is not equivalent to the union of any cells in Figure 4.

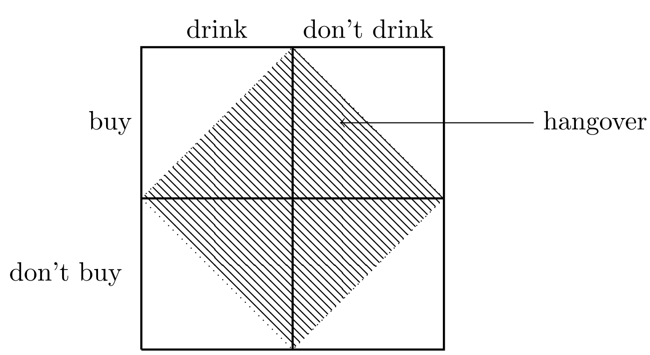

Figure 4: A decision problem that respects the Causal Constraint

By contrast, suppose instead that Carol's decision problem ignored the question of whether to buy alcohol:

$$\{\mathsf{drink}, \mathsf{no\ drink}\}$$

This decision problem violates the Causal Constraint. It represents Carol as deliberating over whether to drink, but not whether to buy libations. This is depicted in Figure 5, where the proposition *Carol buys* is not equivalent to the union of any of the cells of her decision problem.

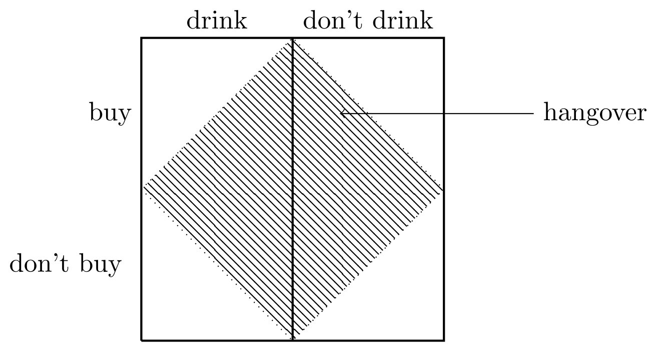

Figure 5: A decision problem that violates the Causal Constraint

To see our full theory in action, consider Figure 6. Here we add doxastic and conative alternatives to the picture in order to arrive at a full model of Carol's mental state. Her doxastic alternatives are blue; her conative alternatives are red. Since her conative alternatives are a subset of her doxastic alternatives, her conative alternatives reflect her beliefs about means-end relations.

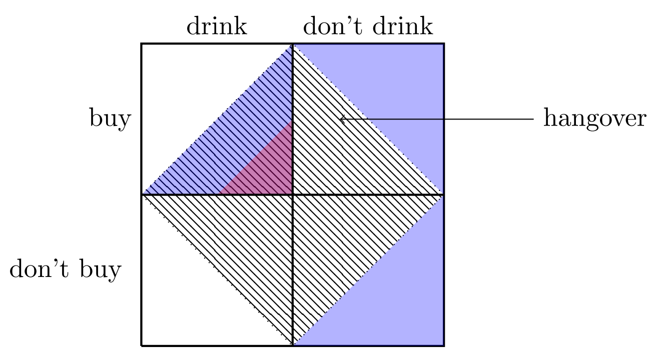

Figure 6: An example

Now consider what this figure says about Carol's intentions. First, she intends to drink, because all of her conative alternatives are worlds where she does so, and drinking is defined on her decision problem. Second, she intends to purchase potations, which she believes is a necessary means to drinking. This is because all of her conative alternatives are worlds where she buys alcohol, and this action is defined on her decision problem. Third, she does not intend to get a hangover. True, in all of her conative alternatives she gets a hangover. But getting a hangover is not defined on her decision problem. Moreover, this asymmetry in definedness is permitted by our Causal Constraint. While drinking is defined on her decision problem, this does not require that its crapulent aftermath is also defined. The latter is a foreseen effect of the former, not a necessary means thereto.[16]

We have motivated the Causal Constraint by appealing to the idea that if you seriously consider whether to bring about p, and you believe that some action q which is under your control is a necessary means of bringing about p, then you are under rational pressure to consider bringing about q. But one might wonder: why doesn't this requirement extends to foreseen effects? Why not hold that if you seriously consider whether to bring about p, and you believe that some effect e which is also under your control will follow from p, that you are under rational pressure to consider performing e?[17]

In response, two points are worth noting. First, our intuitions about cases like Tactical Bomber suggest that there is an asymmetry between believed means and believed consequences of our intended actions: we are under rational pressure to intend the former, but not the latter. Second, there is a principled reason for thinking this asymmetry extends to what we are rationally required to seriously consider. Seriously considering some action is computationally expensive; there are limits to what the possibilities that we can seriously consider at a single time

[16]A few readers have suggested that failures of control trigger a lack of truth value rather than falsity. The idea is that sentences like (3-b) presuppose that William III had control over whether to avoid nuclear war with with France. While we are not fully persuaded by intuition, there are ways to reformulate our analysis to capture this judgment. The simplest reformulation would maintain that A *intends* p is defined only if p is defined on A's decision problem. However, given that presuppositions project out of negation, this would predict that *Carol doesn't intend to have a hangover* is undefined, whereas intuitively this sentence is true. A better approach would be to distinguish the agent's entire decision problem from a second partition, a 'control problem' that divides logical space into cells that reflect outcomes the agent believes they can control. The control problem would refine the decision problem, since the decision problem reflects only some of the actions an agent could take (for example, getting a hangover would be defined on the control problem but not the decision problem). We could then propose that A *intends* p is defined only if p is defined on A's control problem; when defined, A *intends* p is true iff p is defined on A's decision problem and true throughout the cognitive alternatives.

[17]Thanks to a referee for raising this question.

(cf. Bratman 1987). So we must prioritize. There is a clear reason for prioritizing the foreseen means of what we intend; if we don't seriously consider performing those actions, we are unlikely to succeed in achieving our intentions. By contrast, we can achieve our intentions even if we do not seriously consider performing their foreseen consequences. To illustrate with Carol: seriously considering whether to buy alcohol will increase the probability she achieves her goal of drinking, whereas seriously consider whether to get a hangover will not increase the probability she achieves this goal.[18]

We now turn to consider some questions that arise for our theory, as well as some choice points for developing it further. §5 considers possible philosophical interpretations of the two key formal tools in our theory—the agent's decision problem and their conative alternatives. §6 investigates whether our theory generates enough instances of closure.

5 Interpretations

Our semantics relies on two formal tools: decision problems and conative alternatives. While we have given rough glosses on both notions, we have not yet provided a detailed account of either. This was by design. Our primary goal in this paper was to solve our two closure puzzles for intention. As we saw in §4, we could achieve this goal without taking a stand on substantive—and inevitably controversial—questions about the analysis of decision problems and conative alternatives. Here the situation is analogous to a Hintikka semantics for knowledge or belief: even absent a reductive account of the doxastic or epistemic alternatives, it provides a fruitful framework for exploring the logics of knowledge and belief.

Still, by supplementing our formal framework with a substantive account of decision problems and conative alternatives, we can glean further insight into the nature of intention. In this section, we consider some choice points for developing a substantive philosophical interpretation of these notions.

[18] In case the reader is not persuaded by this argument, we refer them to §5, where we sketch an alternative procedure for deriving a principled distinction between the foreseen means and the foreseen consequences.

5.1 Interpreting Decision Problems

When we introduced the notion of a decision problem, we said that it comprises various actions that the agent is deliberating about whether to perform. This is a necessary condition on when an action appears in the agent's decision problem. This necessary condition is all we needed to solve our two puzzles (§4). To recap: this condition explains why one cannot intend outcomes that one lacks the conceptual resources to entertain; after all, one cannot deliberate about an option if one lacks the resources to contemplate it. It also explains why one cannot intend outcomes that one believes to be out of one's control; after all, it is not rational to deliberate about something if you are sure you have no control over it. Finally, this necessary condition, when combined with Belief Containment, sufficed to validate Means-Ends Coherence, given the plausible assumption that if one is deliberating about whether to bring about p, one is rationally required to deliberate over actions that one believes to be necessary means to bringing about p.

So even without providing necessary and sufficient conditions on when an option features into a decision problem, we can achieve our explanatory goals in this paper.[19] Still, it is natural to wonder whether this necessary condition could also be leveraged into a sufficient condition. Could we simply define a decision problem as the set of maximally specific actions the agent is deliberating over?

However, this suggestion runs into a problem: if this all that a decision problem amounts to, our theory will have trouble handling the full range of cases of Foreseeing Without Intending. Meet the reflective carouser. Not only does she consider whether to drink, she also considers the hangover that would ensue. (Perhaps the contemplation of the hangover gives her pause, but she decides to imbibe anyway.) So if her decision problem includes the maximally specific actions that she is deliberating between, it will look like this:

$$\{\mathsf{buy\ \&\ drink\ \&\ hangover}, \mathsf{no\ buy\ \&\ no\ drink\ \&\ no\ hangover}\}$$

But then the proposition, *I get a hangover* is defined on her decision problem, since this proposition (when intersected with the relevant restriction on logical

[19] This fact is important for warding off a potential objection to our solution to the puzzles, which is that it is stipulative rather than explanatory. In response, it is true that our solution to the two puzzles relied on various constraints on the decision problem, e.g., the Control Constraint and the Causal Constraint. But these constraints were not merely stipulated; rather, they follow from the assumption that every option in the decision problem is an action that the agent is deliberating about whether to perform.

space) is a cell in her decision problem (buy & drink & hangover). And so we predict that she intends to get a hangover. But this seems wrong: intuitively, the reflective carouser does not intend to get a hangover any more than her unreflective counterpart does.[20]

The upshot: not every maximally specific action an agent is considering automatically shows up in their decision problem. But then what further conditions must be met?

The answer will depend on one's philosophical views about how to carve the intending/foreseeing distinction. This is a matter of controversy, and not one that we can hope to resolve here. Instead, we will confine ourselves to the more realistic goal of showing how one *prima facie* attractive account of this distinction can be leveraged into a reductive characterization of decision problems.

Here's a natural diagnosis: the reason why the reflective carouser does not intend to get a hangover is that she does not believe that getting a hangover will cause her desires to be satisfied. By contrast, she does believe that both buying and imbibing will causally contribute to the fulfillment of her desires.[21] To incorporate this diagnosis into our framework, we could filter out all of the options from an agent's decision problem that they believe do not causally contribute to the satisfaction of their all-things-considered desires.

Here's one way to go about this. First, rather than focusing on maximally specific actions, let us zoom out and look at all of the actions that the agent is considering performing, where these actions need not be mutually exclusive or exhaustive. Call these the agent's 'considered acts.' Of course, this raises familiar problems about how to individuate actions. For our purposes, we can appeal to whatever is the most common sense way of characterizing the actions available to the agent. Ask the reflective carouser what she is thinking about doing, and a natural answer might be:

$$\{\textsf{buy, no buy, drink, no drink, hangover, no hangover}\}$$

Next, let the 'preferred considered acts' be whichever considered acts the agent believes to be causally conducive to the satisfaction of their all-things-considered desires. Our reflective carouser believes that drinking is conducive

[20] Similar points also apply to Tactical Bomber, although that case has a more complex causal structure.

[21] See Masek 2010 for a related account of the intending/foreseeing distinction.

to the satisfaction of their all-things-considered desires, and she also believes that buying is causally conducive to drinking. So her preferred considered acts include:

$$\{\textsf{buy, drink}\}$$

Now, this set is not a partition. So if we were to simply identify the agent's decision problem with their preferred considered acts, we would lose many of the structural features that decision problems possessed on our old model. However, for any set of consistent propositions, there is a simple way of inducing a partition. The relevant partition fits logical space (or a suitable restriction of it) into cells whose worlds agree on which propositions in the set are true (see Hawke 2016). Applying this procedure to our reflective carouser's preferred considered acts, we get:

$$\{\textsf{buy \& drink}, \textsf{buy \& no drink}, \textsf{no buy \& drink}, \textsf{no buy \& no drink}\}$$

We can now propose that the resulting partition is the agent's decision problem. Note that this is equivalent to our earlier example of a decision problem that respects the Causal Constraint (§4.2). As we noted there, the proposition, *I get a hangover* is not defined on this partition. And so we achieve the desired result that the reflective carouser does not intend to get a hangover.

One might wonder whether the resulting proposal still preserves the verdict that the reflective carouser intends to drink. After all, one might worry, the reflective carouser is aware that drinking will lead to a hangover, which she certainly does not desire. So why is drinking included as one of her preferred considered acts? The answer is that the reflective carouser's drinking desire outweighs her hangover aversion: she wants to revel more than she wants to avoid a hangover. So, all things considered, she desires to drink. Given this, the procedure outlined above predicts that the proposition, *I drink* is defined on her decision problem. But since getting a hangover is an effect of the satisfaction of her desires rather than a cause thereof, our procedure predicts that *I get a hangover* is not defined on her decision problem.[22]

The foregoing is offered as a 'proof of concept': we started with an attractive diagnosis of the intending/foreseeing distinction, and showed how it could be leveraged into a reductive account of decision problems—an account that delivers

[22]Of course, we could imagine a different case where the carouser's aversion to a hangover is stronger than her desire to drink. But in this revised case, the claim that the reflective carouser intends to drink is no longer intuitively true.

the right verdict about the reflective carouser. However, for our purposes we need not commit to this particular procedure. A variety of other accounts of the intending/foreseeing distinction could, in principle, be deployed to a similar end.

5.2 Interpreting the Conative Alternatives

Turn next to the philosophical interpretation of the conative alternatives. By now the philosophy of action literature contains a number of promising theories of the functional role of intention. Each of these theories could be adapted to yield an account of the conative alternatives.

First, consider belief-desire theories, according to which intention boils down to some combination of belief and desire (Audi 1973; Davis 1984; Ridge 1998; Sinhababu 2013). A simple way of implementing this idea in our own framework would be to define the conative alternatives as whichever doxastic alternatives are ranked highest by the agent's desires or preferences. This would yield the following version of our theory:

Belief-Desire Version A intends p iff (i) p holds throughout A's most preferred doxastic alternatives, (ii) p is defined on A's decision problem.

Adopting this substantive interpretation of the conative alternatives does not render the definedness constraint on intention redundant. Without the definedness constraint, we would be left with an unadorned Hintikka semantics of the sort discussed in §3. And so we would be back at square one when it comes to solving our closure puzzles. By contrast, the signal contribution of our definedness constraint is to explain such closure failures.[23]

Next, consider cognitivist theories of intention, according to which intention is just a particular sort of belief about what one will do (Velleman 1989; Marušić and Schwenkler 2018). In order to integrate this view into our framework, we could make the bold move of simply identifying the conative alternatives with the doxastic alternatives, yielding:

[23] In the context of a belief-desire theory, our definedness condition can also be used to capture the idea that idea that intention is a tool for non-ideal agents, who lack the cognitive sophistication required to always take into account every decision-relevant fact (cf. Bratman 1987). The idea is that an agent may have very fine-grained preferences, some of which the agent needs to bracket for the purposes of making efficient decisions. For example, suppose that you have decided to bring wine to the party, but have not yet settled on a bottle. Suppose that your most preferred doxastic alternatives are worlds where you bring a specific vintage—say, a 2015 Beaujolais—but you have not yet given this matter any thought. Our definedness condition correctly predicts that you do not yet intend to bring a 2015 Beaujolais, since this option does not yet feature in your decision problem.

Cognitivist Version A intends p iff (i) A believes p, (ii) p is defined on A's decision problem.

Here too, the definedness constraint is vital to the viability of this theory. We've already noted that belief is not sufficient for intention, since one might believe p while thinking one has no control over p ((8-a)$\not\models$((8-b)). So the definedness constraint provides a new tool for defusing objections that arise for a simple cognitivist view.[24]

A third school of thought, championed by Bratman 1987, takes intention to be a *sui generis* propositional attitude. According to Bratman, intention has certain distinctive dispositional features that distinguish it from any amalgam of belief and desire. Our framework is also compatible with this view. To integrate the two, we could reframe the dispositional features identified by Bratman as necessary conditions on the conative alternatives. To give just one example, Bratman argues that intending p involves a disposition to avoid reconsidering whether to bring about p, whereas merely desiring p does not entail any such disposition (1987: 18-19). We could reframe Bratman's observation as a condition on conative alternatives: if $Con(w) \subseteq p$, then the agent is disposed to avoid reconsidering whether to try to bring about p.

For our purposes, we need not choose between these three options. The important point is that our framework is compatible with any number of theories of the functional role of intention.[25] Moreover, our key semantic contribution—the definedness constraint—can be used to overcome some pressing objections to extant theories. This highlights the fruitful interactions between our semantic framework and treatments of intention in the philosophy of action literature.

[24]Cf. Marušić and Schwenkler 2018, who discuss some counterexamples to a simple cognitivist view. They propose to avoid these counterexamples by identifying an intention with a belief about what one will do which is held on the basis of practical reasoning. Our framework encodes a similar idea, since the decision problems are themselves characterized in terms of practical reasoning.

[25]That said, it is worth noting one advantage of the first two options over the third: both of the first two immediately validate Belief Containment, which we needed to underwrite Means-Ends Coherence. By contrast, it is an open question whether the third option validates Belief Containment; here much depends on the details of the dispositional view.

6 Enough Closure?

6.1 Closure Lost

In §2 we argued that any adequate solution to our puzzle will give us enough closure, without giving us too much. Now, we've shown how our theory accounts for attested closure failures. However, some might worry it only achieves this result at the expense of undergenerating intuitive instances of closure.

Recall that we wanted to explain the oddity of:

(1) ??Ana intends to ski at Zermatt and Chamonix this winter, but she doesn't intend to ski at Chamonix this winter.

(2) ??Ana intends to ski at Zermatt and Chamonix this winter, but she doesn't intend ski at either Zermatt or Chamonix this winter.

These data suggest that any plausible semantics for intention should validate the following minimal instances of closure:

Conjunction Distribution If a rational agent A intends $p \wedge q$, then A intends p, and A also intends q.[26]

Conjunction Weakening If a rational agent A intends $p \wedge q$, then A also intends $p \vee q$.[27]

However, nothing in our semantics thus far guarantees that Conjunction Distribution and Conjunction Weakening hold. Consider the decision problem depicted in Figure 7. Here, the agent's decision problem partitions logical space

[26] For further support for Conjunction Distribution, it is helpful to contrast intention reports with *ought* claims. Consider the case of Professor Procrastinate (Jackson and Pargetter 1986). Our professor is invited to review a book. He is the most qualified reviewer, but he is also terribly unreliable. Some have the judgment that (i-a) is true while (i-b) is false:

(i) a. Professor Procrastinate ought to accept and write.$\not\models$
b. Professor Procrastinate ought to accept.

Here, there is a striking disanalogy with intention reports:

(ii) a. Professor Procrastinate intends to accept and write.$\models$
b. Professor Procrastinate intends to accept.

Any situation where (ii-a) is true is also one where (ii-b) is true, unless our professor suffers from irrationality in addition to tardiness. (For further discussion of Professor Procrastinate in relation to a question-sensitive treatment of deontic modals, see Cariani 2013.)

[27] Cf. Hawke 2016, who argues that any plausible restricted closure principle for knowledge will validate the analogous principles involving knowledge.

into two cells: $p \wedge q$, and $\neg(p \wedge q)$. The conjunction $p \wedge q$ is defined on this decision problem, since it is a union of cells in the problem. By contrast, p, q, and $p \vee q$ are all undefined. So if this is a rationally permissible decision problem—and nothing in our theory so far says that it is not—then the agent can intend $p \wedge q$ without intending p, intending q, or intending $p \vee q$.

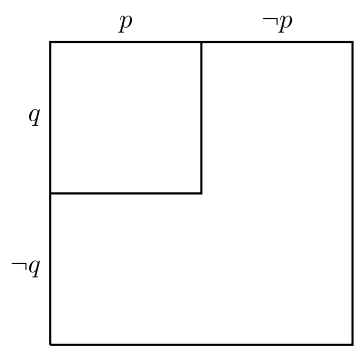

Figure 7: Conjunctive closure failures

6.2 Closure Regained

While this is an important challenge, we think it can be addressed by further unpacking our question-sensitive machinery. Thus far, we have developed our question-sensitive theory of intention in isolation from other semantic and philosophical applications of questions. One major application of questions in the recent literature has been to model *subject matters*. By helping ourselves to developments on this front, we can provide an elegant and independently motivated response to the undergeneration concern.

Some background: researchers who work on subject matters typically assume that every sentence can be assigned to its own subject matter or topic.[28] Here we model this with a topic function $[\cdot]$—a function from sentences to subject matters—defined recursively. For simplicity, following Lewis 1988b, we take the subject matter of an atomic sentence α to be the polar question of *whether* α. So, for example, the subject matter of *It is raining* can be modeled with the polar question, *Is it raining?*, which is in turn modeled as a partition of logical space into the worlds where it is raining and the worlds where it isn't. (Plebani and Spolaore 2021 argue persuasively that atomic sentences are often associated with

[28]See e.g., Yablo 2014, Fine 2015, Hawke 2016, 2018, Berto 2018.

subject matters that are more complex than the corresponding polar question. Instead, they propose that atomic sentences are often focus sensitive, and can be associated with questions that vary either the subject or predicate of the sentence. Everything we say below would carry over smoothly to this proposal.)

On to negation. As noted by Perry 1989 and Hawke 2018, the subject matter of *It's not raining* is intuitively the same as the subject matter of *It's raining*. Hence negation preserves subject matter.

Next, for the cases that are of most interest for us: conjunction and disjunction. Intuitively, the conjunction, *It is raining and it is cold* is about both of its conjuncts. That is, it is partly about the question, *Is it raining?*, and also partly about the question, *Is it cold?*. Likewise, with the disjunction, *It is cold or it is raining*.[29] Summarizing:

Subject Matters Where a is a logically simple sentence:

1. $[a] = \{[\![a]\!], W - [\![a]\!]\}$
2. $[\neg a] = [a]$
3. $[p \wedge q] = [p \vee q] = \{\mathbf{p} \cap \mathbf{q} \mid \mathbf{p} \in [\![p]\!], \mathbf{q} \in [\![q]\!], \mathbf{p} \cap \mathbf{q} \neq \emptyset\}$

For illustration, consider Figure 8, which depicts the subject matters of p, q, and $p \wedge q$. Our account of subject matters guarantees that the subject matter of $p \wedge q$ is at least as fine as that of p and of q.

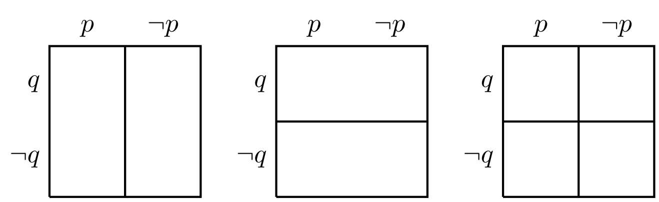

Figure 8: $[p]$, $[q]$, and $[p \wedge q]$

We can generalize from this example by introducing a parthood relation $\leq$ between subject matters. Following Lewis 1988b, we say $[p]$ is part of $[q]$ just in case $[q]$ refines $[p]$, which obtains provided every cell of $[p]$ is a union of cells from $[q]$.[30]

[29] For similar ideas, see Fine 2015; Hawke 2016, 2018; Hoek forthcoming.

[30] When decision problems are restricted to a subset of logical space R, then this definition of parthood can be generalized so that $[p] \leq_R [q]$ iff $\forall \mathbf{p} \in [p] : \exists Q \subseteq [q] : \mathbf{p} \cap R = (\bigcup Q) \cap R$.

Subject Matter Parthood $[p] \leq [q]$ iff $\forall \mathbf{p} \in [p] : \exists Q \subseteq [q] : \mathbf{p} = \bigcup Q$

Now that we are equipped with an account of subject matter parthood, we can put it to work in our semantics for intention. In particular, we can refine our account of when a proposition is defined on an agent's decision problem to reflect subject matters:

Definedness Redefined A proposition p is defined on an agent's decision problem D iff $[p] \leq D$.[31]

According to this definition, definedness is closed under parthood. When q is defined on an agent's decision problem and the subject matter of p is part of the subject matter of q, it follows that p is defined on the agent's decision problem. This definition agrees with our earlier one for atomic sentences, but imposes stronger definedness constraints on logically complex sentences. By doing so, it allows us to validate Conjunction Distribution and Conjunction Weakening.

To see this, note that the theory of subject matter above implies that $[p \vee q] = [p \wedge q]$, and that $[p]$ and $[q]$ are each parts of $[p \wedge q]$. After all, the subject matter of $p \wedge q$ is as fine-grained as the subject matters of p and of q: any cell of $[p]$ is a union of cells from $[p \wedge q]$. So whenever $p \wedge q$ is defined on an agent's decision problem, so are p, q, and $p \vee q$. As a result, whenever someone intends $p \wedge q$, they also intend p, they also intend q, and they also intend $p \vee q$.

Redefining definedness in this way yields further predictions about the logic of intention. For example, it predicts that intentions agglomerate:

Agglomeration If a rational agent A intends p and A also intends q, then A intends $p \wedge q$.

After all, the subject matter of a conjunction is composed of the subject matters of its conjuncts. So if p and q are both defined on the agent's decision problem, so too will be their conjunction.

This prediction seems to be borne out by intuition. Consider:

(9) ?? Ana intends to ski at Zermatt this winter. She also intends to ski at Chamonix this winter. But she doesn't intend to ski at both this winter.

It is hard to make sense of this sort of ascription unless we take Ana to be incoherent (Condoravdi and Lauer 2016; Grano 2017). This should bolster

[31] Cf. Yalcin 2011 on belief and Berto 2018 on imagination.

confidence that our semantics is strong enough to underwrite a wide range of intuitive instances of closure.[32]

Let's take stock. By associating sentences with subject matters, we can refine our account of definedness so as to validate a range of plausible closure principles for intention. Is this solution *ad hoc*? We think not, for three reasons. First, the way we associated sentences with subject matters is independently supported by intuitions about what these sentences are about. Second, our use of subject matters is quite natural, given the prominent role that questions play in both our theory of intention and the standard treatment of subject matters. Finally, the problem of securing enough logical closure is not specific to intention. For example, both knowledge and rational belief go in for Conjunction Distribution and Conjunction Weakening. So we are deploying tools that already earn their keep underwriting restricted closure principles for other attitudes.[33]

[32] While we take it to be a virtue that our account validates Agglomeration, we acknowledge that there are potential grounds for pushback. For example, Goldstein 2016 suggests that Agglomeration gives rise to conative analogues of the preface paradox (Makinson 1965). Suppose Ana intends to ski at Zermatt this winter; ditto for Chamonix, Courchevel, and a dozen other resorts. But she also recognizes that she often tries to pack too much into a single ski season, which usually leaves her exhausted and broke. According to Goldstein, for each of the fifteen ski resorts, she intends to ski there, but she does not intend to ski at all of them. Goldstein's proposed solution is to adopt a threshold semantics for intention that invalidates Agglomeration. The key idea is that intention comes in degrees, and that A intends p just in case A intends p to a sufficiently high degree. Defenders of this threshold semantics could use the question-sensitive machinery developed above by adding the further condition that p is defined on the agent's decision problem. For further discussion of the idea that intention comes in degrees, see Shpall 2016; Dellsén and Sharadin 2017; Beddor 2020.

[33] Some may wonder whether our proposed solution still undergenerates instances of closure. In particular, a referee raises the following data:

(i) a. ? Ana intends to buy a dog, but Ana doesn't intend to buy a mammal.
b. ? Ana intends to buy a dog, but Ana doesn't intend to buy a four-legged creature.

The referee suggests that these reports are infelicitous, which provide evidence that intending to buy a dog entails intending to buy a mammal, and also entails intending to buy a four-legged creature. But nothing in our semantics so far validates these entailments.

In response, we should note there is a delicate question about the data here. While we agree that (i-a) is apt to sound odd out of the blue, it sounds perfectly fine if Ana doesn't believe that all dogs are mammals. Moreover, even if Ana does have this belief, we can arguably rig up a context in which (i-a) is acceptable. Imagine that Ana believes that dogs are mammals, but does not care a whit about their status as mammals; if they belonged to some other biological class, this would suit her just as well. Suppose she intends to buy a dog; consequently, she purchases an adorable beagle from her local petshop. In the evening, she is surprised—but not saddened—to learn that biologists have re-classified dogs as a special type of furry bird. Does it follow that some of Ana's morning intentions were frustrated? This is not obvious. (Similar points can be made using (i-b).)

For those who do not share these judgments, a more concessive reply is also possible. There is nothing to stop us from imposing further requirements on an agent's decision problem in particular cases. In particular, we could require that whenever Ana's decision problem is defined on buying a dog, it is also defined on buying a mammal and on buying a four-legged

7 Comparing Intention and Desire

Our theory starts with a Hintikkan treatment of intention and enriches this framework with question-sensitivity. However, some might wonder, why be beholden to the Hintikkan framework at all? After all, some of the leading semantics for desire ascriptions depart radically from the Hintikkan mold. Perhaps these analyses provide a promising template for analyzing intention.

In this section, we argue against this idea. We start by showing that desire and intention differ systematically in their closure profiles. Consequently, the leading semantics for desire would make the wrong predictions about the closure profile of intention—a point we make using Heim's seminal 1992 semantics as a test case.

7.1 Closure Puzzles for Desire

One major motivation for rejecting a Hintikka semantics for desire is that desire is not closed under logical consequence. However, when we look more closely, we find that the closure failures for desire depart in systematic ways from the closure failures for intention.

Start with the Puzzle of Logical Consequence. We saw that one class of counterexamples to Unrestricted Logical Closure came from the control requirement on intention. However, desire is not subject to any such control requirement: we can desire outcomes that we know we cannot influence. Suppose Jason is watching the latest football game on television, certain that nothing he can do will influence the outcome. Whereas (10) is false, (11) is intuitively true:

(10) Jason intends for the Jaguars to win.

(11) Jason wants the Jaguars to win.

There are also counterexamples to a logical closure requirement on desire that do not arise for intention. Consider an example discussed by Asher (1987); Heim (1992); and Levinson (2003), among others. Nicholas would rather not pay to charter a private plane, but would be happy to fly private if a wealthy friend offered to let him fly for free. In this case, (12-a) is true while (12-b) is false:

creature (in particular, by being defined on the question that answers both whether she buys a mammal and whether she buys a four-legged creature). We also wish to leave it an open question whether there are further more general requirements on an agent's decision problem—requirements that could be folded into the framework developed here.

(12) a. Nicholas wants a free trip on a private jet. $\not\models$
b. Nicholas wants a trip on a private jet.

In this respect, desire patterns differently from intention:

(13) a. Nicholas intends to take a free trip on a private jet. $\models$
b. Nicholas intends to take a trip on a private jet.

On the most natural way of filling out the scenario, both (13-a) and (13-b) are intuitively false. Of course, we could add details to make (13-a) come out true—for example, if Nicholas has started cajoling his jet-owning pals for a free flight. But these details render (13-b) true as well.

So desire gives rise to a version of the Puzzle of Logical Consequence. But the puzzle is different from that which arises for intention. This gives reason to doubt that these two puzzles will receive a uniform solution.

Does the Puzzle of Believed Consequence also arise for desire? This is open for debate. It does seem that an analogue of Foreseeing Without Intending applies to desire: the Tactical Bomber may desire to destroy the enemy munitions factory without desiring the foreseen consequence of civilian causalties. What about Means-Ends Coherence? On the one hand, there is some appeal to the idea that agents are under rational pressure to desire the believed means to their desired ends. On the other hand, the intuitive force of this idea seems much weaker than in the case of intention. Moreover, a number of theorists have proposed counterexamples to a Means-Ends Coherence requirement on desire. Phillips-Brown 2018 describes a case where Al knows he can only go to a concert if he takes a long drive that will likely make him car-sick. In this case, Phillips-Brown argues that both (14-a) and (14-b) seem true:

(14) a. Al wants to see the concert.
b. Al doesn't want to take the long drive.

Similarly, Villalta 2008 suggests that the following inference is invalid:

(15) a. I want to teach Tuesdays and Thursdays next semester.
b. I believe that I will teach Tuesdays and Thursdays next semester if and only if I work hard. $\not\models$
c. I want to work hard now.

If we reject a Means-Ends Coherence requirement on desire, then desire does not give rise to a Puzzle of Believed Consequence at all. But then we should not expect the correct theory of desire reports to carry over to intention reports.

So we've seen that there are reasons to expect the analyses of desire and intention to diverge. As a consequence of this, existing analyses of desire struggle to predict the closure profile we've ascribed to intention. For reasons of space, we make this point by focusing on one particular analysis, due to Heim 1992. We focus on Heim's account because it is among the most prominent analyses of desire in literature, and because Heim herself suggests that this semantics can be used as a template for analyzing intention. However, many of the points we make apply equally well to other semantics for desire, such as Levinson 2003's decision-theoretic analysis.

7.2 Heim's Comparative Analysis of Desire

Building on Stalnaker 1984, Heim 1992 develops a counterfactual semantics for desire reports. An agent wants p just in case she believes that if p were to happen, things would be better than if $\neg p$ were to happen. More precisely, for any two worlds w and v, say $w > v$ just in case w is more desirable than v to the agent (A). Then, for propositions p and q, say that $p > q$ iff every world in p is more desirable than every world in q. Then let $Sim_w(p)$ be the worlds most similar to w where p obtains. Heim proposes:

$$[\![A \text{ wants } p]\!]^w = 1 \text{ iff } \forall v \in Dox(w) : Sim_w([\![p]\!] \cap Dox(w)) > Sim_v(W - [\![p]\!] \cap Dox(w))$$

Heim's semantics is particularly relevant for our purposes, because she briefly discusses an extension of her semantics to intention. Heim suggests that intention is similar to desire, except that it has different domain of quantification. Instead of quantifying over the agent's doxastic alternatives, intention reports quantify over a wider set of worlds: those that are compatible with the agent's beliefs "about matters unaffected by [their] own future actions" (1992: 199). Where $Dox^-(w)$ denotes this set:

$$[\![A \text{ intends } p]\!]^w = 1 \text{ iff } \forall v \in Dox^-(w) : Sim_v([\![p]\!] \cap Dox^-(w)) > Sim_v(W - [\![p]\!] \cap Dox^-(w))$$

To evaluate this proposal, consider how it applies to our puzzles. An initial observation is that it invalidates Unrestricted Logical Closure for both desire and intention. This is one of Heim's main goals: she wants to explain why Nicholas' desire for a free trip on a private jet does not entail a desire for a trip on a private jet. To see how the semantics delivers this result, note that Nicholas prefers the closest worlds where he gets a free trip to the closest worlds where he does not. But this does not imply that he wants a trip on a private jet, since the closest worlds where he takes such a trip are worlds where he pays for it. While this is a desirable result when it comes to desire attributions, it seems incorrect in the case of intention, since the correspond inference involving intention is valid ((13-a)$\models$(13-b)). Relatedly, Heim's semantics does not explain why intention, but not desire, is subject to a control requirement. It thus doesn't explain the difference in truth-value between (10) and (11).

Turn next to the Puzzle of Believed Consequence. Heim's semantics does not validates Means-End Coherence for either desire or intention. Recall Barry, who both intends and desires to cook a delicious feast, and also believes that buying groceries is a necessary means to do so. Does he desire to buy groceries? To answer this, Heim's semantics instructs us to take an arbitrary doxastic alternative w, and find the closest doxastic alternative v where he buys groceries. We then compare v's desirability with that of the closest world u where he does not buy groceries. But here is the problem. The closest world where he buys groceries may not be a feast-making world. After all, a believed necessary means is not always a believed sufficient means; it might be that in some of the doxastic alternatives where Barry buys groceries, he botches his culinary preparations. If v is one of these worlds, then it will not be more desirable than u. (Better to save money on the groceries if the meal will be a disaster!)

Heim's semantics also has trouble accommodating the full range of cases in which one desires or intends an outcome without desiring or intending its foreseen consequences. In particular, Heim's theory validates the following restricted closure principle:

Closure Under Believed Material Equivalence If a rational agent A believes that p iff q, then A desires or intends p iff A desires or intends q.

To see this, suppose the Tactical Bomber believes the biconditional: *civilians will die if and only if the factory is bombed.* And suppose they desire to bomb the factory. Then for an arbitrary doxastic alternative w, the closest doxastic alternative where the factory is bombed (v) is preferred to the closest doxastic alternative where it is left undisturbed (u). But since every doxastic alternative where the factory is bombed is a world where civilians are killed, it follows that the closest doxastic alternative where civilians are killed is preferred to the closest alternative where they live (since $v > u$). So by Heim's semantics, Tactical Bomber wants the civilians to be killed. The point extends straightforwardly to intention. After all, the Tactical Bomber may believe the relevant biconditional is true regardless of how the Tactical Bomber acts. (Regardless of whether the Tactical Bomber is the one who bombs the factory, either both claims will be true or both will be false.)[34]

8 Conclusion

The last several decades have seen a booming industry of increasingly nuanced analyses of propositional attitude reports. One theme in this literature is that not all propositional attitudes conform neatly to the Hintikkan mold. However, much of this literature has focused on desire ascriptions, with less attention to intention. At the same time, a rich parallel tradition in philosophy of action has made great strides understanding the nature of intention; however, it has paid little heed to the semantic details of intention ascriptions in natural language.

This paper has sought to bridge these traditions by developing a new semantics for intention reports. The account developed here belongs to a broader research program that takes propositional attitudes to be question-sensitive. However, our account deployed this question-sensitive machinery in novel ways, using it to explain some of the puzzling closure properties that distinguish intention from other propositional attitudes.

While our theory is primarily a contribution to the analysis of intention, it has implications for future applications of question-sensitivity. Most of the question-sensitive accounts of propositional attitudes in the literature have

[34]While we have focused on Heim's view, similar problems arise if we attempt to translate other leading semantics for desire into a semantics of intention. For example, Levinson 2003 defends a decision-theoretic account of desire, on which an agent desires p when the expected utility of p is greater than of $\neg p$. The resulting theory invalidates Unrestricted Logical Closure in the free trip example. In addition, this theory invalidates Means-End Coherence, while validating Closure Under Believed Material Equivalence.

For another example, Phillips-Brown 2018 puts a question-sensitive spin on a Hintikka semantics for desire, in the tradition of von Fintel 1999. On von Fintel's view, *A believes p* is true iff p holds in all of the most preferred of A's doxastic alternatives. Phillips-Brown's key innovation is to make the desirability of a world relative to a question the agent is considering. One consequence of this semantics is that it invalidates Means-Ends Coherence for desire. While this may be a good consequence when it comes to desire reports, this is a non-starter as a theory of rational intention. Phillips-Brown also predicts that an agent can desire p relative to one question and desire $\neg p$ relative to another, because each question generates its own desirability ordering. By contrast, our framework does not allow questions to flip the desirability ordering (reflected by the conative alternatives), hence we do not make this prediction.

not imposed any constraints on the relevant questions. By contrast, we have shown that by imposing substantive constraints on practical questions, we can validate specific rational requirements on intention. In this regard, our approach structurally resembles the way that a Hintikka semantics models rational and logical properties of attitudes reports using specific constraints on accessibility relations. We suggest that researchers working in a question-sensitive setting should follow suit; a fruitful avenue for future research is to examine which constraints on questions will validate different rational requirements on different propositional attitudes.[35]

[35]Thanks to two anonymous referees, as well as Kyle Blumberg, Milo Phillips-Brown, Sam Carter, Dan Hoek, Ben Holguín, Cameron Domenico Kirk-Giannini, and participants in a workshop on Intention, Action, and Language for helpful feedback. We are also grateful for the support of a Tier 1 research grant from the Singapore Ministry of Education on the topic, 'Epistemology Beyond Belief' (R-106-000-062-115).

References

Nicholas Asher. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy*, 10:125–197, 1987.

Robert Audi. Intending. *Journal of Philosophy*, 70(13):387–403, 1973.

Bob Beddor. Fallibility for Expressivists. *Australasian Journal of Philosophy*, 98 (4):763–777, 2020.

Jonathan Bennett. Morality and consequences. In *The Tanner Lectures on Human Values*. University of Utah Press, 1980.

Francesco Berto. Aboutness in imagination. *Philosophical Studies*, 175:1871 – 1886, 2018.

Francesco Berto. Simple hyperintensional belief revision. *Erkenntnis*, 84:559 – 575, 2019.

Michael Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

John Broome. Normative requirements. *Ratio*, 12:398–419, 1999.

Fabrizio Cariani. 'ought' and resolution semantics. *Noûs*, 47(3):534–558, 2013.

Cleo Condoravdi and Sven Lauer. Anankastic conditionals are just conditionals. *Semantics and Pragmatics*, 9:1–61, 2016.

Mark Davies. The Corpus of Contemporary American English (COCA): 520 million words, 1990-present, 2008-. URL https://corpus.byu.edu/coca/.

Wayne Davis. A Causal Theory of Intending. *American Philosophical Quarterly*, 21(1):43–54, 1984.

Finnur Dellsén and Nathaniel Sharadin. The Beliefs and Intentions of Buridian's Ass. *Journal of the American Philosophical Association*, 3(2):209–226, 2017.

Fred Dretske. Epistemic operators. *Journal of Philosophy*, 67:1007–1023, 1970.

Daniel Drucker. The attitudes we can have. *The Philosophical Review*, 129(4): 591–642, 2020.

Luca Ferrero. Can I only intend my own actions? In Shoemaker, editor, *Oxford Studies in Agency and Responsibility*, volume 1, pages 70–94. Oxford University Press, Oxford, 2013.

Kit Fine. Angellic content. *Journal of Philosophical Logic*, 2015.

Simon Goldstein. A Preface Paradox for Intention. *Philosophers' Imprint*, 16 (14):1–20, 2016.

Thomas Grano. The logic of intention reports. *Journal of Semantics*, 34:587–632, 2017.

Jeroen Groenendijk and Martin Stokhof. Studies on the semantics of questions and the pragmatics of answers, 1984.

Charles Hamblin. Questions. *Australasian Journal of Philosophy*, 36(3):159–68, 1958.

Daniel Harris. Imperative Inference and Practical Rationality. *Philosophical Studies*, 179:1065–1090, 2022.

Peter Hawke. Questions, topics, and restricted closure. *Philosophical Studies*, 173, 2016.

Peter Hawke. Theories of aboutness. *Australasian Journal of Philosophy*, 96(4), 2018.

John Hawthorne. *Knowlege and Lotteries*. Oxford University Press, New York, 2004.

Irene Heim. Presupposition Projection and the Semantics of Attitude Verbs. *Journal of Semantics*, 9:183–221, 1992.

Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.

Daniel Hoek. Minimal rationality and the web of questions. In Andy Egan, Dirk Kindermann, and Peter van Elswyk, editors, *Unstructured Content*. Oxford University Press, forthcoming.

Ben Holguín. Thinking, believing, and guessing. *Philosophers' Imprint*, 22, 2022.

Frank Jackson and Robert Pargetter. Oughts, options, and actualism. *The Philosophical Review*, 95(2):233–255, 1986.

Joshua Knobe. Intentional action and side effects in ordinary language. *Analysis*, 63:190 – 193, 2003.

Niko Kolodny and John Brunero. Instrumental rationality. *The Stanford Encyclopedia of Philosophy*, 2020.

Angelika Kratzer. The Notional Category of Modality. In Eikmeyer and Rieser, editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. W. de Gruyter, Berlin, 1981.

Dmitri Levinson. Probabilistic model-theoretic semantics for 'want'. *Proceedings of SALT*, 13, 2003.

David Lewis. Relevant implication. *Theoria*, 54(3):161–174, 1988a.

David Lewis. Statements partly about observation. *Philosophical Papers*, 17(1): 1–31, 1988b.

David Makinson. The paradox of the preface. *Analysis*, 25(6):205–207, 1965.

Matt Mandelkern, Ginger Schultheis, and David Boylan. Agentive modals. *Philosophical Review*, 126:301 – 343, 2017.

Berislav Marušić and John Schwenkler. Intending is Believing: A Defense of Strong Cognitivism. *Analytic Philosophy*, 59(3):309–340, 2018.

Lawrence Masek. Intentions, motives, and the doctrine of double effect. *Philosophical Quarterly*, 60(240):567–585, 2010.

Alison McIntyre. Doctrine of double effect. *The Stanford Encyclopedia of Philosophy*, 2019.

Jeff McMahan. Intention, permissibility, terrorism, and war. *Philosophical Perspectives*, 23:345–372, 2009.

Dana Kay Nelkin and Samuel C. Rickless. Three cheers for double effect. *Philosophy and Phenomenological Research*, 89(1):125–158, 2014.

Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, 1981.

John Perry. Possible worlds and subject matter. In *The Problem of the Essential Indexical and other Essays*. CSLI Publications, Palo Alto, CA, 1989.

Milo Phillips-Brown. I want to, but... *Proceedings of Sinn und Bedeutung*, 21(2): 951–968, 2018.

Matteo Plebani and Giuseppe Spolaore. Subject matter: A modest proposal. *Philosophical Quarterly*, 71(3):605–622, 2021. doi: 10.1093/pq/pqaa054.

Warren Quinn. Actions, intentions, and consequences: the doctrine of double effect. *Philosophy and Public Affairs*, 18(4):334–351, 1989.

Michael Ridge. Humean Intentions. *American Philosophical Quarterly*, 35: 157–178, 1998.

Craige Roberts. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5(6):1–69, 1996/2012. doi: 10.3765/sp.5.6.

Jonathan Schaffer. Contrastive knowledge. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 1, pages 235–271. Oxford University Press, Oxford, 2005.

Jonathan Schaffer. Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403, 2007.

Sam Shpall. The Calendar Paradox. *Philosophical Studies*, 173:801–825, 2016.

Neil Sinhababu. The Desire-Belief Account of Intention Explains Everything. *Noûs*, 47(4):680–696, 2013.

Robert Stalnaker. *Inquiry*. MIT Press, Cambridge, MA, 1984.

Victor Tadros. Wrongful intentions without closeness. *Philosophy and Public Affairs*, 43(1):52–74, 2015.

Alejandro Pérez Carballo. Structuring logical space. *Philosophy and Phenomenological Research*, 2016.

David Velleman. *Practical Reflection*. CSLI Publications, 1989.

Elisabeth Villalta. Mood and Gradability: An Investigation of the Subjunctive Mood in Spanish. *Linguistics and Philosophy*, 31(4):467–522, 2008.

Kai von Fintel. NPI licensing, strawson entailment, and context dependency. *Journal of Semantics*, 162(2):97–148, 1999.

Stephen Yablo. *Aboutness*. Princeton University Press, NJ, 2014.

Seth Yalcin. Nonfactualism about Epistemic Modality. In Egan and Weatherson, editors, *Epistemic Modality*. Oxford University Press, Oxford, 2011.

Seth Yalcin. Belief as Question-Sensitive. *Philosophy and Phenomenological Research*, 97(1), 2018.