

6

Conditional Reasoning Part II

Indicatives, Subjunctives, and the Ramsey Test

6.1 Introduction

Last chapter I focused on the material conditional and its place within systems of natural deduction. I argued that nothing in the evaluation of such conditionals gives good reason to posit *sui generis* imaginative states. Even when experimental prompts abstract away from context to encourage participants to reason purely deductively, participants' conclusions concerning conditionals systematically deviate from the norms appropriate to the material conditional. It would be surprising, then, if systems of natural deduction were nevertheless good models for our actual psychological states and inferential processes. At the beginning of last chapter, I also noted that the indicative and subjunctive conditionals used in everyday life are evaluated differently than the truth-functional material conditional. The way we arrive at beliefs in such conditionals demands a different account, whatever we say about the material conditional. This chapter considers what that account should look like, and whether there is good reason to think that *sui generis* imaginative states will have a role within it.

The Ramsey test looms large in contemporary discussions of indicative and subjunctive conditionals and, as we will see below, inspires some of the best-developed and most influential theories of the role of imagination in conditional reasoning (Nichols & Stich, 2000; Currie & Ravenscroft, 2002; Williamson, 2016). I will explain what the Ramsey test consists in presently (section 6.2). A key argument of this chapter is that, regardless of whether the Ramsey test records an important insight about the nature of conditionals, it gives us no reason to posit *sui generis* imaginative states. Indeed, even if the Ramsey test provides a cogent analysis of how we assess conditionals, a view where such reasoning only involves beliefs is still preferable.

6.2 The Ramsey Test and Its Psychology

Often, when faced with whether to accept 'if p then q ,' we simply consider the likelihood of q , on the supposition that p . This is to calculate the *conditional probability* of q , given p . If q appears likely in the event that p , we are inclined to

believe that if p then q ; if not, then not. If it's true that our degree of belief in a conditional is simply a function of the probability we assign to q , in the event that p should occur, it's no wonder that analyzing indicative and subjunctive conditionals on a par with the material conditional creates "paradoxes" of the kind discussed last chapter. For if indicatives and subjunctives had the same truth conditions as the material conditional, we should be *equally* interested in the probability of not- p when considering whether to accept 'if p then q ' as we are in the conditional probability of q given p . That is, our confidence in the truth of 'if p then q ' should raise as a function of our confidence in 'not- p '. Yet, in general, it does not. As we saw last chapter, Edgington's increasing confidence that the Queen is not home doesn't make her increasingly confident that, if the Queen is home, then she is worried about Edgington's whereabouts. It does, however, sound plausible to say that Edgington's inclination to believe the (indicative) conditional, 'if the Queen is home, then she is worried about my whereabouts,' raises as a function of how likely Edgington takes it to be that the queen will be worried about Edgington's whereabouts, in the event that the Queen is home.

In the case of subjunctive conditionals, where we typically believe the antecedent to be false, the point is even clearer. Our degree of belief in the conditional does not match what it ought to be if the conditional had the truth conditions of a material conditional. After all, we may reject a subjunctive conditional—'if I'd had a healthy breakfast, I would have won the presidential election'—despite knowing that its antecedent is false. All material conditionals with false antecedents are true, however. As with indicative conditionals, it may seem, instead, that we will accept a subjunctive counterfactual to the extent that we find the consequent likely, had the antecedent occurred (though we will see counterexamples to this shortly).

Philosophers and psychologists have thus explored an alternative theory of both indicative and subjunctive conditionals, tracing to Frank Ramsey. In a brief and now very famous footnote, Ramsey offered the following characterization of what occurs when we consider whether to accept a conditional:

If two people are arguing 'If A will C?' and are both in doubt as to A, they are adding A hypothetically to their stock of knowledge and arguing on that basis about C... We can say they are fixing their degrees of belief in C given A.

(1929, p. 143)

Ramsey's proposal is that, when deciding whether to accept 'If A then C,' we are deciding how strongly we ought to believe that C, in the event that A—we are "fixing [our] degrees of belief in C given A." Stalnaker (1968) cites Ramsey's idea in developing his own influential theory of conditionals, applying it to both indicatives and subjunctives. (For reasons we will see, this lumping of the two under a

single analysis is relatively uncommon, with the Ramsey test more commonly applied only to indicatives.) Stalnaker characterizes the psychological process of assessing a conditional as follows:

First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally consider whether or not the consequent is then true. (1968, p. 102)

Often, when we evaluate a conditional (such as a counterfactual), we already *disbelieve* the antecedent. In such cases, Stalnaker adds, “you cannot simply add it to your stock of beliefs without introducing a contradiction” (1968, p. 102). This is why he goes beyond Ramsey in specifying that “adjustments” must be made to one’s stock of beliefs to accommodate the antecedent. Of course, in such cases, we do not literally come to *believe* the antecedent during hypothetical reasoning. Nor is it obvious what it could mean to add a proposition to one’s stock of beliefs *only hypothetically*. After all, *something* really needs to be done in carrying out the Ramsey test, if it is to be seen as a bit of reasoning by which we decide whether, and how strongly, to believe a conditional. We *really* need to add the antecedent to our stock of beliefs *only hypothetically*. What could this involve?

Several philosophers see imagination as the crucial ingredient. Currie and Ravenscroft (2002) (“C&R”) suggest that there is “nothing new” in their suggestion that imaginative states are belief-like because “philosophers interested in belief dynamics assume that imagining is belief like... when they offer the Ramsey test as a way of deciding whether you should accept a conditional” (2002, p. 12). In cases where we don’t believe that *P*, but wish to determine whether we should accept ‘If *P* then *Q*’ we can simply add *P* “in imagination... since imagination preserves the inferential patterns of belief” (p. 12). This is what C&R think Ramsey and his followers have in mind when they speak of disputants “adding *A* hypothetically to their stock of knowledge.” If *Q* then “emerges as reasonable” one knows that ‘If *P* then *Q*’ should be believed (p. 12).

In a similar spirit, Nichols & Stich (2000) (“N&S”) hold that, in order to form beliefs in conditionals relevant to pretending that *p*—e.g., “If we were at a tea party, then pastries would be served”—one must first represent that *p* in the Possible Worlds Box (hereafter, the “PWB,” and aka the “Imagination Box”) and see what inferences come to be made on that basis “in” the PWB. It is then by importing the latter representations—e.g., *q*₁ & *q*₂ & *q*₃...—into our Belief Box as consequents of a conditional with *p* as its antecedent that we arrive at beliefs of the form: “If it were the case that *p*, then it would be the case that *q*₁&*q*₂...&*q*_{*n*}” (2000, p. 128). As with C&R, the idea is that, in order to find out what would happen if *p*, we need to represent (in the PWB, or with *sui generis* belief-like

imaginative states more generally) that p and see what we infer “in imagination.” This is, putatively, a way of finding out what one *would* infer, *were* one to in fact believe that p ; it is a way of adding p to one’s stock of beliefs “only hypothetically.” And it is available to us just because the relevant form of imagination is belief-like in its inferential properties. In N&S’s terms, imagination’s being belief-like amounts to the same “inference mechanisms” operating on representations in the PWB as in the Belief Box—a feat made possible by the representations in each box occurring “in the same code.”

Finally, while Williamson (2016) does not place the same emphasis on imaginative representations having a certain format or “code,” his account of imagination’s role in conditional reasoning is also inspired by Ramsey’s observation that “how we evaluate conditionals is closely tied to how we update our beliefs on new information” (2016, p. 118).

6.3 From Belief Conditions to Truth Conditions

Before evaluating these proposals concerning the psychological implementation of the Ramsey test, we should pause to consider the relationship of the Ramsey test to theories of conditionals more generally. As noted, Stalnaker endorses the Ramsey test as an account of what we do when *considering* a conditional (both indicative and subjunctive)—an account of the psychological conditions under which we will *believe* a conditional. The question still to be answered for a theory of conditionals, he notes, is how we are to “make the transition from belief conditions to truth conditions; that is, to find a set of truth conditions for statements having conditional form which explains why we use the method we do use to evaluate them” (1968, p. 102). To that end, he appeals to the notion of a *possible world*, which is to serve as an “ontological analogue” to a hypothetical set of beliefs. Stalnaker’s idea is that, for any hypothetical set of (consistent) beliefs, there is a possible world where those beliefs are all true. He then defines a conditional’s truth conditions in terms of such possible worlds—applying it to both indicatives and subjunctives—as follows:

Consider a possible world in which A is true, and which otherwise differs minimally from the actual world. *‘If A , then B ’ is true (false) just in case B is true (false) in that possible world.* (p. 102)

By this criterion, the truth or falsity of ‘If A then B ’ doesn’t hang at all on the actual truth value of A . Instead, it hangs on whether, in whichever possible world where A is true (i.e., in whichever “ A world”) that is otherwise most similar to the actual world, B is true as well. (Lewis (1973) develops a similar and equally influential approach to the truth conditions of counterfactuals.) An

immediate advantage is that not all counterfactual conditionals come out true. “If Jeremy had asked Alice to dance, she would have agreed” is true just in case, in the possible world most similar to our own where Jeremy asks Alice to dance, she agrees.

How are we to *know* whether B is true in the A world most similar to our own? Here we are brought back to the Ramsey test. Adding A “hypothetically” to our set of beliefs and adjusting the set minimally for consistency with A constitutes an effort toward mentally representing a minimally different possible world where A is true. If we judge B to be likely from within that hypothetical set of beliefs, our confidence in the truth of ‘If A, then B’ should rise correspondingly. Such is Stalnaker’s influential suggestion for linking the Ramsey-inspired epistemology to the metaphysics of conditionals. Notably, this theory still allows us to *misjudge* the truth value of a conditional. The A world that is *in fact* most similar to our own can easily differ from the possible world corresponding to a particular person’s idiosyncratic belief set when it is revised to accommodate A.¹

6.4 A Difference for Subjunctives

But there is a more profound way in which gaps can open between the Ramsey test and a conditional’s truth value—one that has importance to theories of imagination. Earlier I mentioned that philosophers often approach the truth conditions of indicative and subjunctive conditionals differently, with Stalnaker (1968) being an exception. Why think that the mood difference between past-tense subjunctive conditionals (“If he had asked Sally to dance, she would have accepted”) and past-tense indicative conditionals (“If he asked Sally to dance, then she accepted”) warrants any difference in the treatment of their epistemology or metaphysics? A series of examples from Adams (1970) is telling.

¹ Bill Lycan (2001) offers an example where the two diverge: “If I finish this chapter today, Norway will have an unusually early autumn in 2055.” The Ramsey test can explain why he does not believe this conditional, even if he suspects he will not finish the chapter today. For when he considers the belief set that results from adding “I will finish this chapter today” to his own stock of beliefs and revises for consistency, the proposition that Norway has an early autumn in 2055 does not emerge as likely. However, Lycan notes:

Suppose that unbeknownst to us and even the world’s most competent physicists, there are arcane laws of nature *L* such that the conjunction of *L* with my finishing this chapter today entails Norway’s having the early autumn... Stalnaker or Lewis would count the conditional as straightforwardly true, since a world in which I finish the chapter but Norway fails to have the early autumn would have to differ from our world in its laws of nature (a large difference). (2001, p. 70)

On both Stalnaker (1968) and Lewis’s (1973) theories, the truth of a conditional hangs on what is true at the “closest” or “most similar” possible world where the antecedent is true; and, as Lycan notes, differences in the natural laws governing a world are generally held to be among the most significant differences there can be. In this case, our denial of the conditional after applying the Ramsey test is “just a case of perfectly well justified false belief.”

Consider this famous pair of conditionals that differ only in mood:

- (1) If Oswald didn't kill Kennedy, then someone else did.
- (2) If Oswald hadn't killed Kennedy, then someone else would have.

Most judge (1) to be true and (2) false. The difference in mood seems to generate a difference in truth value. Examples are easily multiplied. (It seems obvious that, if a meteor did not lead to the dinosaurs' demise, then something else did. It is less obvious that, if a meteor hadn't led to the dinosaurs' demise, then something else would have.) The recognition that subjunctives and indicatives have different truth conditions is the departure point for much contemporary theorizing about conditionals. Like any orthodox view, it can be questioned (see, e.g., Edgington 2008; Lassiter, 2017). But I won't do so here; my aim instead is to explain the relevance of this common view to theories of imagination.

My focus will be on the different epistemologies pertaining to each kind of conditional. Our tendency to accept (1) could reasonably be thought to derive from an application of the Ramsey test. We add 'Oswald didn't kill Kennedy' hypothetically to our stock of beliefs and make the minimal adjustments needed to maintain consistency (*viz.*, we remove the belief that Oswald killed Kennedy). We then consider whether 'someone else killed Kennedy' should also be a member of that set—that is, we determine how *probable* that proposition is, from within the terms of our hypothetical stock of beliefs. The probability will be high. For still within that set remains the proposition that Kennedy was killed. Any *minimal* adjustment to one's set of beliefs—made only to preserve consistency—must leave that one in place. And indeed it is crucial to the relevance and reliability of the Ramsey test that our adjustments *be* minimal. Practically any consequent could appear to follow with probability from the truth of any antecedent, given *arbitrary* adjustments to a surrounding stock of beliefs, after all. A version of the Ramsey test that allowed for *more than* minimal adjustments to one's belief set would not be a reliable method for arriving at the truth of (1).

N&S—who are more explicit than C&R and Williamson in their explanation of how, precisely, we implement something like the Ramsey test—hold that when considering whether to accept 'if *p* then *q*,' we represent *p* in imagination and combine it there with representations of *all the other contents we believe*,² making only minimal adjustments for consistency. They posit an "UpDater" mechanism that is responsible for these adjustments—one that works either by deleting incompatible representations from the PWB, or by filtering them out before they arrive in the PWB. This gives us, in imagination, a set of representations consistent with a conditional's (not believed) antecedent that is otherwise coextensive

² "Let us assume that in addition to the pretense initiating premise, the cognitive system puts the entire contents of the Belief Box into the Possible Worlds Box" (Nichols & Stich, 2000, p. 123).

with our beliefs. We then let the relevant inferences unfold “in” the PWB (2000, pp. 122–5). In the present case, one such inference will be: someone other than Oswald killed Kennedy. We will then come to believe that, if Oswald didn’t kill Kennedy, then someone else did.

Subjunctive conditionals present a problem for this method, however. Recall (2): ‘If Oswald hadn’t killed Kennedy, then someone else would have.’ When I add ‘Oswald didn’t kill Kennedy’ hypothetically to my set of beliefs and revise only for consistency, it will still emerge as reasonable that someone else did. For within that stock of beliefs remains the belief that Kennedy was shot. Lest I succumb to conspiracy theories, I would be wrong to infer that, if Oswald hadn’t killed Kennedy, someone else would have. Nor am I inclined to make the inference.

For these reasons, it is thought by many that the Ramsey test is not the method we in fact use when considering subjunctive conditionals.³ If that is so, then theories of imagination that aim to explain conditional reasoning in Ramseyan terms—showing how general principles governing the updating of beliefs apply *mutatis mutandis* to *sui generis* imaginative states—cannot extend themselves to explaining our reasoning about subjunctives. This is a significant limitation on these approaches. Imagination is a dear friend of the subjunctive mood. One of the key platitudes surrounding ‘imagining,’ in the A-imagining sense, is that it is a kind of cognition involved in considering how things *could have been*, had something else not occurred. Philosophers, in particular, are interested in the role imagination plays in our understanding of counterfactuals, linking them to our understanding of causation and natural laws. If we analyze imagination as a faculty which, by default, moves forward inferentially from a proposition as one would if one believed the proposition, imagination becomes ill-suited to explain counterfactual reasoning. Yet this appears to be precisely what theorists such as Nichols & Stich, Currie & Ravenscroft, and Williamson have in mind when they propose that “left to itself, the imagination develops the [possible] scenario in a reality-oriented way, by default” (Williamson, 2016, p. 116).

This does not mean that there is no possible explanation to be had for reasoning about subjunctive conditionals from within the terms of such accounts—only that they will require significant amendment to explain how belief-like *sui generis* imaginative states are used in reasoning about subjunctive conditionals. It has been proposed, for instance, that subjunctive counterfactuals simply require us to “rewind time” back to when the antecedent occurs, revising all beliefs concerning matters subsequent to that time. We then “re-run” the tape forward from there, as it were, assessing whether the consequent is probable in *that* context

³ Problems have also been noted in extending the Ramseyan approach to indicatives, though these appear less pervasive. Counterexamples trade on cases where the truth of indicative conditional’s consequent would require the agent considering the conditional not to know its truth. An example of this attributed by van Frassen (1980, p. 503) to Richard Thompson is: “If my business partner is cheating me, I will never realize that he is” (Bennett, 2003, pp. 28–9).

(Edgington, 2008). This is a version of what Lassiter (2017) calls, the “Rewind, Revise, Re-run” heuristic: “Rewind to the antecedent time, Revise to make the antecedent true, and selectively Regenerate following events that depend causally on the antecedent” (p. 527). This is one way of preserving a “suppositional” account of subjunctive counterfactuals—one where we “assume true” a certain set of propositions and assess the probability of another (the consequent) from within the terms of that set.

Note, however, that on this view we are no longer able to make use of our existing belief set minus one or two, together with our ordinary inferential procedures taken “offline.” Instead we need to mentally rewind world history to a certain point—knowing what to edit out from our beliefs in the process and what to let stand—and then let it play forward again, calculating the consequences of the antecedent’s having been true. Any heavy-duty suppositional account, like that of Nichols & Stich (2000), will need to posit additional features of cognitive architecture to accommodate these time-relative adjustments to one’s belief set. Thus, evaluating subjunctive conditionals will require not simply the “offline” reuse of ordinary inference procedures—whatever they may be—but some mechanism capable of determining the historical time represented by each of our beliefs. This is like having to posit a special mechanism for weeding out all and only beliefs about events occurring in Europe, or events occurring during the summer. There is nothing incoherent in the idea of such a mechanism. But its addition, as a posit, is a significant cost to the theory.

I will analyze the Ramsey test in more depth soon when considering indicative conditionals. Those deeper reflections, aimed at undermining the assumed link between the Ramsey test and *sui generis* imaginative states generally, will apply *mutatis mutandis* to subjunctive conditionals. Supposing, until then, that we wanted to avoid a Ramseyan/suppositional account of subjunctive conditionals, how *else* might we understand the psychology of evaluating subjunctives? Let ‘ $\square \rightarrow$ ’ stand for the if-then relation of subjunctive conditionals. On the Stalnaker/Lewis account of subjunctive conditionals, which remains the most influential, to judge ‘ $p \square \rightarrow q$ ’ true is (roughly) to judge that the closest (i.e., most similar to our own) possible worlds where p are also worlds where q . The question then becomes how we go about judging the relative location of possible worlds with respect to our own. How do we decide that all possible worlds where Oswald doesn’t shoot Kennedy and Kennedy is still assassinated are further from the actual world than those where Oswald doesn’t shoot Kennedy and no one else does?

On the one hand, we do not have a complete mystery here. Most will agree that a world where the laws of nature are the same as our own is closer to our own than one where the laws are very different (Lewis, 1973). And specific cases seem easy to judge: a world where Donald Trump was never elected president seems closer than any where a centipede was elected president (unless, of course,

Trump's not winning would have required an unobvious adjustment in natural laws not required by a centipede's victory...). Proceeding from uncontroversial cases, we can make efforts to systematize those nearness judgments, showing how certain principles (such as sameness of natural laws) drive many of them. In the end we seem destined to arrive at a set of norms, heuristics, and rules of thumb for making relative location judgments. Without question, this is a complex form of reasoning that depends, crucially, on what we take to be true in the actual world. Our ignorance about the actual laws of nature will spill over into ignorance about the relative closeness of other possible worlds. But there is no reason to think that engaging in this sort of complex reasoning would require us to set aside, or step outside of, our beliefs. The process is *fueled by* belief. Based on what we believe true of our own world—and what we take “closeness” to depend upon—we reason that a possible world where A and B are true is closer to our own than any where A and not-B is true, and so judge that that “ $A \square \rightarrow B$ ” is true. It was only the Ramsey test—with its metaphorical talk of adding a proposition hypothetically to an existing stock of beliefs—that gave the impression of our needing to somehow step outside of our beliefs through the use of *sui generis* imaginative states to engage in this complex reasoning.

6.5 The Ramsey Test and the Psychology of Indicative Conditionals

Whatever difficulties subjunctive conditionals may present to theories of imagination modelled on the Ramsey test, such views would remain attractive if they still offered the best account for our reasoning with indicatives. These accounts—specifically, those of N&S, C&R, and Williamson—all posit *sui generis* imaginative states. Their success at explaining our reasoning with indicative conditionals would give reason to doubt my claim that we can fruitfully explain the A-imagining at work in conditional reasoning entirely in terms of beliefs. We need to look more closely now at how such theories propose to explain indicative conditional reasoning, to assess whether that reasoning could instead be implemented entirely within one's beliefs.

To begin, it is important to see that one can endorse the Ramsey test as an adequate account of the epistemology of indicative conditionals without thereby committing to the existence of *sui generis* imaginative states. The Ramsey “test”—with its language of adding propositions hypothetically to one's stock of beliefs—can be seen as an artful way of expressing the *theory* that, when evaluating an indicative conditional, people calculate the probability of the conditional's consequent given the antecedent. Taken as a theory about how people evaluate indicatives (or *should* evaluate them), it is at odds with the theory that people evaluate

indicatives as though they are truth-functional in the manner of the material conditional. The latter theory predicts that people will evaluate indicatives by consulting the truth or falsity of their antecedents and consequents and matching those values to the corresponding row of the truth table for the material conditional. This debate about what people are actually up to when evaluating indicatives can march forward without either side committing to a view about *how* people go about judging conditional probabilities (when they do).⁴

And so it does in the book *If*, by psychologists Jonathan St. B. T. Evans and David Over (2004). Evans & Over hold that the Ramsey test captures the psychology of indicative conditionals to the extent that people's judgments concerning the probability of a conditional are shown, in experimental settings, to mirror their judgments about the conditional probability of the consequent, given the antecedent.⁵ "The majority of people judge the probability of 'if p then q' to be at least close to the conditional probability of $P(q|p)$," they explain. This is "precisely what we would expect if people were conforming to the Ramsey test" (2004, p. 154). Yet they also stress that "the Ramsey test does not tell us *how* people make conditional probability judgments," where answering *how* would be to provide an account of the psychological mechanisms or processes by which the inference is made (p. 169, emphasis added). They add:

Trying to answer the question of how the Ramsey test is implemented is a problem... there is no one answer to the question. There are rather many answers that will refer to many psychological processes. (p. 25)

And again:

The Ramsey test is a high-level description of many processes that contribute to hypothetical thought. Describing fully the processes that can make up a Ramsey test of a conditional is a formidable challenge for psychological research on conditionals and in judgement and decision making. (p. 158)

We can see N&S, C&R, and Williamson as making specific (and indeed ambitious) proposals for how the reasoning behind the Ramsey test occurs—proposals that invokes *sui generis* (or "offline") imaginative states. They are not merely

⁴ It is possible to hold that, while people in fact treat indicative conditionals as though their probability is equivalent to the conditional probability of the consequent given the antecedent, such conditionals nevertheless have the truth conditions of the material conditional. This would require holding that the proper normative account of indicative conditionals does not mesh well with how indicatives are in fact treated by ordinary reasoners.

⁵ Johnson-Laird & Byrne (2002) adduce their own set of experiments where participants do not accept or reject conditionals based on the probability of the consequent, given the antecedent, to argue that the Ramsey test gives an incorrect account of our actual psychological engagement with conditionals.

concurring with Ramsey, and Evans & Over, that indicative conditionals are evaluated differently than material conditionals.

I want now to explore whether some approach along those lines proposed by N&S, C&R, and Williamson *must* be true. Having answered in the negative, I'll then evaluate whether such proposals are likely simpler or more powerful than any that appeal exclusively to belief. Again I will reply in the negative. "Belief-only" approaches are in fact more parsimonious, as they require nothing over and above what theories that posit *sui generis* imaginings themselves require—while doing without the *sui generis* imaginings.

6.6 A General Argument Against the Need for *Sui Generis* Imaginative States in Conditional Reasoning

Before we can find a belief-only approach to conditional reasoning plausible, we need to see how one is possible. That is the project of this section. Judgments in favor of conditionals are, for those espousing the Ramsey test, judgments concerning the high probability of the conditional's consequent, given its antecedent. Is it possible to judge, and thereby come to know, that the conditional probability of q given p is high, using only one's standing beliefs? In at least some cases, the answer is obviously yes. Evans and Over note that "there must always be some conditionals, of the form 'if p then q ', that people consider probable *to the extent that they have learned that q type events follow p type events*" (2004, p. 8, emphasis added). If I believe that every time I have gone to the grocery store there were flowers for sale, it seems I can reasonably infer from my beliefs alone that if I go to the grocery store, I can buy flowers there. The inference would be something along the lines of:

1. Whenever I've gone to the grocery store, flowers have been for sale.
2. There is no reason to think conditions have changed.
3. Therefore, the probability that I can buy flowers, in the event that I go to the grocery store, is high.
4. Therefore, if I go to the grocery store, I can buy flowers there.

I did not, in conducting this reasoning, need to imagine that *I am at the grocery store* and see what emerged as likely. Williamson, N&S, and C&R will likely agree that simple inferences of this kind can be made without use of a PWB or an "offline" imaginative exercise. The more difficult question is whether there are nevertheless *some* more complex judgements in favor of conditionals that could not be carried out in this "belief-only" way. Evans & Over seem to think that there are, noting that people are not restricted to reasoning of the sort just described, "but can also use, for example, mental models of complex causal relationships to

make probability judgments about conditionals...probability judgments that are generated theoretically and not just on the basis of past experience” (2004, p. 9).

We already looked closely, last chapter, at the notion of a “mental model” as it appears in the work of Johnson-Laird and colleagues. Properly understood, mental models are constituents of ordinary beliefs (or so I argued). Here I want to focus on whether there is any reason, in principle, to think that *some* judgments in favor of conditionals simply couldn’t be construed as justifiably flowing from a set of preexisting beliefs—and if, instead, there are some that require the use of *sui generis* imaginative states. As a means to arguing that there is no such barrier on belief-only approaches, I will defend the following entailment:

RT Entailment: If an application of the Ramsey test, via *sui generis* imaginative states, would lead *S* to justifiably infer ‘if *p* then *q*,’ then *S* already has, prior to triggering the *sui generis* imaginative states, beliefs on the basis of which she can justifiably infer ‘if *p* then *q*.’

To deny the RT Entailment, one needs to hold that an application of the Ramsey test, through the use of *sui generis* imaginative states, would lead *S* to reasonably infer ‘if *p* then *q*,’ despite the fact that *S* did not, prior to conducting the Ramsey test, have beliefs that warranted inferring “if *p* then *q*.” (That would be to affirm the antecedent of the RT Entailment while denying its consequent.) The problem here is that the objector would *also* need to hold that there were no additional beliefs *S* needed to gain, so as to be warranted in inferring ‘if *p* then *q*.’ After all, the only thing that ensures that a person’s Ramsey test will return *q* as likely after *p* is added hypothetically to her stock of beliefs “in imagination” is the stock of beliefs itself: do they, or do they not, contain other information that, when combined with *p*, make *q* emerge as likely? So, prior to imagining that *p* in the service of any Ramsey test that results in one’s judging that *if p then q*, one’s stock of beliefs must already be such that, were *p* imagined, *q* would emerge as likely. If that is not true of one’s stock of beliefs, then imagining that *p* will not cause *q* to emerge as likely, and ‘if *p* then *q*’ will not be inferred.

Thus, the person denying the RT Entailment is in the absurd position of holding that one can lack justification for inferring ‘if *p* then *q*,’ despite there being no further beliefs one needs to gain before justifiably inferring ‘if *p* then *q*.’ What else *does* the person need to do before justifiably inferring ‘if *p* then *q*?’ The RT Entailment-denier will propose that *S* has to carry out a very specific psychological procedure of adding *p* hypothetically to her beliefs—via an offline imaginative exercise, use of the PWB, *sui generis* imaginative state, or similar—before justifiably inferring ‘if *p* then *q*.’ That claim is both adventurous and question-begging. In every other situation where a person has all the beliefs she needs to justifiably infer a proposition, we do not hesitate to conclude that she can use those beliefs—without aid of some other type of content-bearing mental state—to arrive at the inference. We may not yet be able to articulate the specific

inferential rule, heuristic, or process that would be used in each case. But neither has the defender of *sui generis* imaginative states articulated such for imagination—appealing, as they do, only to “belief-like” or “reality-oriented” inference patterns. To the extent that questions remain open about the principles governing such inferences, they apply to theories invoking imaginative states as well.

So, the RT Entailment is hard to deny. If the Ramsey test rightly captures a normative standard for when we are justified in inferring a conditional, then *sui generis* imaginative states are never needed to provide justification for those beliefs. Nevertheless, it may be that such states provide a necessary *tool* for inferring conditionals—that they are not normatively but, rather, *psychologically necessary* for arriving at (at least some of) our beliefs in conditionals. This is the idea I want to explore now.

6.7 It Is Simpler to Just Use Beliefs—Considering an Example from Williamson

We have seen that there is no normative barrier to doing without *sui generis* imaginative states when carrying out the mental calculations needed to reason in accord with the Ramsey test. Beliefs will serve just fine, in principle. But it might still seem that such reasoning would be less cognitively demanding, faster, and, indeed, only psychologically feasible were we to use imaginative states that are distinct from our beliefs—that it is some cognitive limitation of our own that prevents us from being able to rely solely upon beliefs in all cases of conditional inference. In developing and responding to this objection, I will look in some detail at an example provided by Williamson (2016) in his explanation of imagination’s role in conditional reasoning. Aside from the fact it is adduced in favor of there being “offline” imaginative states at work in conditional reasoning, there is nothing exceptional about the case. My argument will be that, in cases such as these, there are always background beliefs in generalizations that can serve as sufficient fuel for the inference, without need of *sui generis* imaginings. Further, it is no more cognitively demanding or time-consuming to make use of those background beliefs; for the very same background beliefs are needed on approaches that invoke *sui generis* imaginings as well.

In analogizing conditional reasoning to ordinary, non-conditional inference, Williamson describes a shepherd who is told that the sheep have broken out of their pen. On the basis of that testimony, the shepherd infers that the sheep have gone down to the river. “Presumably,” Williamson writes, “even if the shepherd had not been given the testimony, he could still have reached the indicative conditional conclusion ‘If the sheep have broken out of the pen and disappeared, they have gone down to the river.’” Williamson thinks of the two different inferences—one from testimony, the other only hypothetical—as involving “online” and “offline” processes, respectively:

If we regard the shepherd's updating of his beliefs in the first case [involving testimony] as an online process, then we can regard his evaluation of the conditional in the second case as the corresponding offline process. If he accepted the conditional on the basis of an imaginative exercise . . . then that imaginative exercise is the offline analogue of online updating. Very roughly, the online and offline processes take the same input—"The sheep have broken out of the pen and disappeared"—and deliver the same output—"The sheep have gone down to the river"—by the same means. One process is online and the other offline in virtue of the different sources of the input. (p. 118)⁶

Williamson's proposal is openly inspired by the Ramsey test (2016, p. 118) and similar to N&S's theory of hypothetical reasoning, if less explicit in its details. For Williamson, the imaginative reasoning process takes place "offline," even if it is an "analogue" of "online updating" (p. 118). Like N&S, Williamson hypothesizes that during an imaginative episode "various offline cognitive procedures add further conclusions to a pool that starts with the initial supposition." He adds, importantly, that "most of the procedures are non-deductive" (p. 120). The inferences that, for Williamson, occur "offline" nevertheless mirror the broadly inductive or abductive inferences we would make were we to believe the supposition. "Left to itself," Williamson explains, "the imagination develops the scenario in a reality-oriented way, by default" (p. 116). The idea that imagination has a "default" mode where it develops scenarios in a "reality-oriented" way is again similar to N&S and C&R's claim that imagination is belief-like in preserving the inductive and deductive inference patterns that characterize transitions among beliefs. Unlike N&S and C&R, however, Williamson explicitly allows that this kind of "reality-constrained" imagining can incorporate imagistic states, in addition to language-like states.⁷

⁶ Williamson explains that during most imaginings the cognitive processes will take "a mix of online and offline input." Yet, in the block quotation above, he remarks that "one process is online and the other offline in virtue of the different sources of the input." If the difference in type of input is to constitute the difference between online and offline processes, it is not clear how one of those processes can involve "a mix" of input types (the process itself would then be "mixed" between an online and offline process). What Williamson seems to mean is that, usually, the states used by our offline imaginative processes involve a mix of *contents*, insofar as some of the contents are believed and some are not believed. Understood in that way, he arrives at N&S's view, where the majority of the representations in one's Belief Box are duplicated within imagination, with a few premises inserted that are not also believed. (Subsequently, on N&S's view, the contents that conflict with the new solely-imagined premises are said to be weeded out by the UpDater mechanism.) The presence of all this information "in imagination" is needed to explain how non-deductive, belief-like inferences—inferences entirely shaped by one's contingent beliefs—can unfold in imagination, without imagination being "online." It is also one of the significant *costs* of this approach to conditional reasoning.

⁷ A difficulty in interpreting Williamson on the psychology of conditional reasoning is that, in some of his examples, he seems to favor a view where the knowledge relied upon in counterfactual reasoning is sensory in nature and "may be stored in the form of some analogue mechanism, perhaps embodied in a connectionist network, which the subject cannot articulate in propositional form" (2007, p. 145). (This seems to apply to his stream-jumping case (2016, p. 118), and the

With this general picture in place, we can map out the reasoning that Williamson takes to occur when the shepherd reasons hypothetically about the sheep breaking out of their pen:

Offline Sheep Counting

1. The shepherd registers the question, “Where are the sheep likely to be if they have broken out of their pen and disappeared?” and begins relevant processing.
2. The shepherd represents, “in imagination” (or “in the PWB” or “offline”), that the sheep have broken out of their pen and disappeared.
3. Beliefs about the dispositions of the sheep and their typical behaviors are accessed, and their contents are copied into imagination (or “in the PWB” or “offline”) so as to be “mixed” with the representation that the sheep have broken out of their pen. Such contents include propositions such as that the sheep like to drink water and frolic in the river and that the sheep have, in the past, run down to the river when their pen was left open.
4. It is inferred, “in imagination” (or “in the PWB” or “offline”), that the sheep have gone down to the river. (The principle or process by which this inference is achieved is important to consider—we will return to it.)
5. On the basis of this processing in imagination, the following conditional is inferred and takes up residence in one’s Belief Box: “If the sheep are out of their pen and have disappeared, then they have gone down to the river.”

First, a few notes on step 3. This “mixing” of what is imagined with contents that are believed is necessary for the imaginative episode to develop “in a reality-oriented way” and, especially, for it to develop in ways that mirror what would be one’s inductive inferences were one to believe the antecedent. We saw that N&S, in their explanation of how the mixing occurs, propose that the entire contents of one’s Belief Box are copied into one’s PWB and then revised for consistency with the conditional’s antecedent. With all of that information in the PWB, the same “inference mechanisms” that operate on belief can then operate within the PWB to draw out inductive and abductive inferences that mirror those that would be

rock-rolling-down-the-hill case, from Williamson (2007, p. 143).) Yet, in other examples, he seems to have in mind processing that involves articulable inference rules and propositionally stored knowledge. (This seems to apply to his sheep escaping from the pen case (2016, p. 119), and the case of counterfactual reasoning about who would have won a general election in Britain in 1948, had there been one (2007, p. 150). There can be no clarity on the distinction between the two kinds of reasoning at work in each case until the distinction between “analogue” and “propositional” forms of thought is clearly drawn, which Williams does not do. He holds that each rely on a more general ability “to predict the future” (2007, p. 150). On the face of it, accurately predicting the future requires lots of knowledge about the past; it does not, however, require—or even suggest—*sui generis* imaginative states or “offline” processing.

drawn were one to believe the antecedent. Williamson does not offer details on how the mixing occurs, other than that “the imagination develops the scenario in a reality-oriented way, by default” and takes “a mix of online and offline inputs.”

As we saw when considering the Ramsey test in its application to indicative and subjunctive conditionals, we will not infer offline what we would have inferred from an antecedent, online, unless most of our beliefs are made relevant to the inference.⁸ Thus, as N&S recognize, the amount of information that must be brought to bear in any “offline” inference of the sort Williamson imagines is vast, provided that the imagining is indeed to develop the antecedent in a “reality-oriented way.” Step 3, above, only highlights a fraction of the background beliefs that will be relevant to the inference moving forward.⁹

Now let’s consider the nature of the inference made in step 4, where the shepherd infers, in imagination, that the sheep have gone down to the river. We know that it is supposed to be the same type of inference that the shepherd would have made had he instead come to believe, through testimony, that the sheep had broken out of their pen. Certainly, it is not a *deductive* inference that the sheep have gone down to the river. Rather, like most acts of conditional reasoning, the shepherd’s inference is based on past experience. Williamson agrees that the inferences that take place during “imaginative exercises” usually “depend somehow on past experience, and go beyond it non-deductively” (p. 119). Thus, the shepherd’s reasoning in step 4 appears to be something like an inductive inference, made in light of the information brought to mind in step 3 (though, as we will see, Williamson is not entirely happy calling the inference “inductive”). The issue I want to consider now is whether we can follow Williamson and N&S in understanding the shepherd’s inference as essentially inductive in nature, drawing on beliefs about how things have typically gone in the past, while rejecting their claim that the reasoning makes use of states that are not beliefs. We know, from the arguments above, that such a position is coherent in principle. To assess whether it is *psychologically* plausible, it will help to dig into the details of a specific example.

Consider, first, how we might generate a “light-duty” (see Chapter 2) explanation of the shepherd’s coming to infer the conditional about the escaping sheep. To do so, we just need to think about what the shepherd himself might say when asked to justify his inference. Suppose we ask him, “Where are the sheep likely to be, if they have broken out of their pen and disappeared?” He thinks a moment

⁸ Recall, for instance, the indicative conditional “If Oswald didn’t kill Kennedy, then someone else did,” which most judge to be true. If, in deliberating on whether to accept this conditional, we fail to “mix” the proposition that Kennedy has been killed with the proposition that Oswald didn’t kill Kennedy, we will wrongly judge the conditional to be false.

⁹ One might wonder whether an account such as Williamson’s could make do with only a small sub-set of one’s (relevant) beliefs being copied into imagination, as opposed to most of them. I will address this response below.

and replies: “If they have broken out of their pen and disappeared, then they have gone down to the river.” “Why do you think that?” we ask. “Well, the sheep love going down to the river for a drink and a frolic,” he replies. “That’s where they tend to go when out of their pen and left to their own devices.” From these expressed beliefs alone—with nothing occurring “offline”—he can reasonably infer that if the sheep have broken out of their pen, then they have gone down to the river. Who are we to gainsay him for doing so?

The shepherd is not revealing to us the specific computational process he has exploited, or making any heavy-duty claims about the nature of beliefs themselves. His is a superficial, light-duty folk psychological explanation for how and why he came to infer the conditional that he did. But it is significant in itself. If we were inclined to describe him as having imagined his way to the answer, we now see that this case of imagining can alternatively be described as a case of drawing inferences from his standing beliefs in past regularities and tendencies. Further, as we saw in Chapter 2, superficial folk psychological explanations are often suggestive of a heavy-duty counterpart. That is again the case here. We can think of the shepherd’s prior beliefs as mental representations that are causally implicated in the generation of a new belief in a conditional. Those beliefs are activated, so as to take part in a chain of reasoning, when the shepherd’s cognitive system registers the question of the sheep’s likely whereabouts when out of their pen.¹⁰ With these points in mind, we can map out the psychological processes as follows:

Online Sheep Counting

1. The shepherd registers the question, “Where are the sheep likely to be if they have broken out of their pen and disappeared?” and begins relevant processing.
2. Beliefs about the sheep’s preferences and tendencies are accessed from the Belief Box, such as “The sheep like to drink water and frolic in the river” and “The sheep have, in the past, gone down to the river when their pen was left open.”
3. From these beliefs, the following conditional is inferred and takes up residence in one’s Belief Box: “If the sheep are out of their pen and disappeared, then they have gone down to the river.”

Note, first, that steps 1 and 2 are the same as steps 1 and 3, respectively, in *Offline Sheep Counting*. The episode requires the shepherd to begin by registering a

¹⁰ As discussed earlier (in Chapter 5, fn. 7), simply registering a question, so as to reason about its answer, does not suggest a role for *sui generis* imaginative states. For instance, we do not need to imagine simply in order to ask ourselves “What is the capital of Ohio?” or “What kinds of things are sold at Starbucks?” Whether *answering* these self-put queries requires *sui generis* imaginative states is the point at issue. (See also section 8.8.)

particular question and searching his memory for information relevant to answering. Again a handful of beliefs of particular relevance are accessed. However, in *Online Sheep Counting*, there is no “mixing” of states in imagination (or the PWB), no tokening of a mental representation with the content “The sheep have broken out of the pen and disappeared.” Instead, the two standing beliefs cited in step 2 serve as the mediating states that take us from “Where are the sheep likely to be, if they have broken out of their pen?” to an inference in favor of the conditional “If the sheep are out of their pen and disappeared, then they have gone down to the river.” Again the inference is not deductive in nature, but based on past experience—it is, we can say, *broadly* inductive, so as to leave open the precise inductive inference rule deployed.

The alternative I have just described is certainly simpler than the “offline”/PWB version in one respect: it does not require the interaction of two different kinds of mental states (*sui generis* imaginings, and beliefs); nor, for that matter, does it require the wholesale copying of one’s Belief Box into some other area of the mind, and the subsequent adjustment of that copied set for consistency with the hypothetical antecedent. Supposing both alternatives are live options, why prefer more complex account?

One response may be that we have yet to carefully consider the *inference rules* that govern the transitions among states—either within the Belief Box, in step 3 of my proposed account, or “offline” (or in the PWB), in step 4 of the Williamson/N&S account. If the kinds of rules or heuristics used on the belief-only account were much more complex or difficult to implement than those used on the two-attitude/offline accounts, that would be a point in favor of the latter. But there is no reason to think that the inference rules exploited on the two-attitude/offline accounts would be of a different kind. The key inferences are not deductive on either picture; so, in neither case are there easily stated rules to be imported from systems of natural deduction. In both cases, the inferences are inductive or abductive in nature. On the *sui generis* imaginings and “offline” views, after registering the question, “where have the sheep gone, if they have broken out of their pen?” we move, in imagination, from the propositions that (a) the sheep have broken out of their pen and disappeared, (b) the sheep like to drink water and frolic in the river, and, (c) the sheep have, in the past, gone down to the river when their pen was left open, to the (imagined) conclusion that (d), the sheep have gone down to the river. From there we arrive at the belief (e) that if the sheep have broken out of their pen, then they have gone down to the river. These inferences are inductive, abductive, or probabilistic; it is difficult to extract a schematic inferential principle that is deployed (as is indeed the case with *all* inductive inferences, as Goodman (1983) showed).

On the belief-only view, after registering the question, “where have the sheep likely gone, if they have broken out of their pen?” we move from (a) the sheep like to drink water and frolic in the river, and (b) the sheep have, in the past, gone

down to the river when their pen was left open, to the conclusion that (c) if the sheep have broken out of their pen, then they have gone down to the river. Again the inference is inductive at heart, with the precise schematic inference rule or heuristic that is followed being a matter for empirical investigation. But there is no reason to think that such inferences are easier to carry out “offline” (“in imagination”) than online—and, indeed, no reason to think that the inferences would be of a different sort, were they to occur offline, instead of online. So the only salient difference between *Offline Sheep Counting* and *Online Sheep Counting* is that the offline version builds in states and inferential steps that aren’t necessary in the online version.

The online account is clearly preferable.

Perhaps anticipating this sort of objection, Williamson shows ambivalence when considering whether the kinds of inferences that take place during his “imaginative exercises” are simply inductive in nature:

Could someone argue that what have here been called ‘imaginative exercises’ are really just inductive inferences? Most of them depend somehow on past experience, and go beyond it non-deductively. If that suffices for a cognitive process to be an inductive inference, then they are inductive inferences. (2016, p. 119)

Yet he qualifies this point, asserting that the inferences:

Do not depend on the subject’s *remembering* the relevant past experiences. What matters is whether they have made the subject skillful enough in performing the imaginative exercise itself. It is irrelevant to the process whether the subject can assemble the particular premises of the supposed inductive inference. Nor is it remotely clear in the given cases how to fill in the ‘F’ and ‘G’ in the conclusion of the supposed inductive inference, ‘All Fs are Gs’ (or ‘Most Fs are Gs’, for that matter)... the imaginative exercises are inductive inferences only in a sense so loose as to be entirely unhelpful. (p. 120)

On the one hand, we can agree with Williamson that in calling the relevant inferences inductive, we do not arrive at a deep understanding of the inference rules and principles that underlie the reasoning. It is less clear why Williamson finds it *more* helpful to describe the inferences in terms of a skillful “offline” imaginative exercise—as though those notions are more perspicuous. By definition, the idea of an “offline” inference, or thought-transition, is clear *only* to the extent that we already understand the kinds of inductive or probabilistic inference rules and heuristics that drive the ordinary “online” inferences that the “offline” processes supposedly mirror. Consider again the scenario where, on the basis of testimony, the shepherd comes to *believe* that the sheep have broken out of their pen and disappeared. The shepherd then infers that the sheep have gone down to the river.

This is not a deductive inference, obviously. Here, too, it is unlikely that the shepherd makes use of a formal inductive premise of the form “All observed Fs have been Gs.” Nevertheless, the inference is, in some sense, inductive or probabilistic. Would Williamson find it “entirely unhelpful” to call it inductive, because doing so does not give us a clear picture of the nature of the inference rule deployed? In that case, neither do we have a clear picture of what an imaginative exercise involves, as it must (on his picture) make use of the very same sort of process or inference rule. Invoking the notion of an imaginative exercise won’t move us any further forward than the loose appeal to induction he disparages.

6.8 Mental Imagery and Conformations of the Brain?

The arguments so far considered speak mainly to those who see in the Ramsey test a validation of the idea that conditional reasoning involves use of *sui generis* belief-like imaginative states. Some of what Williamson says about imaginative exercises is suggestive of a less intellectualist picture, where it is not so much a set of belief-like imaginings (copied from one’s actual beliefs) that is relevant to the conditionals one will infer, but rather the way one’s brain has been hard-wired through past experience. Recall that, in questioning the view that imaginative exercises are inductive inferences based on past experiences, Williamson holds that the exercises “do not depend on the subject’s *remembering* the relevant past experiences.” What matters, he tells us, “is whether they have made the subject skillful enough in performing the imaginative exercise itself” (p. 120). Here Williamson seems to have in mind a picture where people are able to generate imaginative states (including sequences of mental images) that represent likely ways a certain scenario would unfold, and where their doing so does not depend upon an ability to exploit their beliefs in a stepwise inference. In earlier work, he offers similar remarks in an explanation of the relation between our beliefs and the counterfactual judgments we are inclined to make:

Very often, the background knowledge needed to evaluate a counterfactual consists not of specific items of information acquired on specific occasions but of a more general sense of how things go, honed over long experience. Such a sense is typically not presented to the subject in usable verbal form...Of course, underlying the inarticulate sense of how things go must be some conformation of the brain, but the latter does not constitute a theory from which the subject can infer the counterfactual or its negation. (2005, p. 14)

Here Williamson is addressing subjunctive counterfactuals as opposed to indicatives; yet he shows a similar inclination to move away from talk of “reasoning”

towards a “conformation of the brain” that allows one the requisite “skill” to arrive at what otherwise look to be paradigmatic *inferences*.

Williamson is welcome to invoke the brain and its conformations here, but doing so does nothing to support the idea that conditional reasoning occurs “offline.” At bottom, the proposal is for a kind of reasoning that is “honed over long experience,” dependent on the shape of the brain, and difficult to put into words. If there is such reasoning, its relation to imagination is obscure. Williamson *may* simply mean to pick out reasoning that involves mental imagery.¹¹ But, in that case, “offline” imaginative exercises are simply I-imaginings, as I have defined them. And we have already seen (in Chapters 3 and 4) that there is no conflict in I-imaginings reducing to more basic folk psychological states.

Note that to remove beliefs from one’s picture of conditional reasoning—opting instead for a notion of “imaginative skill”—is to remove whatever explanatory value the Ramsey test may have had in shedding light on the nature of that reasoning. For the Ramsey test *does* explicitly appeal to the notion of a belief set, and of adding propositions to it hypothetically. Once it is divorced from the notion of belief, the posit of an “offline” imaginative process can no longer be explicated by appeal to the online inferential procedures it supposedly mirrors. Inferences hold among beliefs, after all, not perceptual experiences; and perceptual experiences, presumably, would be the online counterparts to “offline” mental imagery, on Williamson’s picture.¹² It is possible to propose something *like* purely imagistic inference rules, appealing to mechanisms involved in sensorimotor planning and prediction (Langland-Hassan, 2016; Van Leeuwen, 2011). One could even propose that it is through the use of imagistic states that we compute the conditional probability of a consequent given an antecedent. But this raises again the question of whether the imagistic states mediating those predictive computations are themselves properly viewed as judgments, *sui generis* imaginings, or something else.

Alternatively, one could reject the Ramsey test with its suggestion that, in evaluating an indicative conditional, we calculate the conditional probability of the consequent given the antecedent. In that case, an alternative theory of the computation undertaken when evaluating a conditional is required before we can move forward with an account of the particular mental states and inference rules used in carrying it out. After all, if we don’t know *what* we are trying to determine when evaluating a conditional, we certainly can’t know whether our doing so requires states other than beliefs.

¹¹ Though we also saw, in Chapter 3, that Williamson warns us not to “over-generalize to the conclusion that all imagining involves imagery” (2016, p. 117).

¹² Not surprisingly, when commenting on the nature of the inference rules that are used during imaginative exercises, Williamson moves back to a belief-inspired picture, according to which, “the deductive aspect of the whole process will look something like the method of tableaux in first-order logic, by which the consequences of the initial premises are teased out” (2016, p. 120).

6.9 Two Objections Considered

I will end this chapter by considering two objections to the line of argument I've so far developed. The first asks whether views positing *sui generis* imaginings must really be committed to the claim that, when we imagine, most of what we believe is “copied into” imagination. The second holds that there are specific acts of counterfactual reasoning—philosophical thought experiments being prime examples—where we infer a new generalization, and where our doing so requires focusing on the details of a hypothetical case. This may be thought to create difficulties for the kind of “belief-only” account I've sketched above, which relies upon our having pre-existing background beliefs in relevant generalizations. I will address these objections in turn.

6.9.1 Would We Really Have to Copy So Much into Imagination?

The first objection holds that I have overstated the requirements of theories invoking “offline” states or *sui generis* imaginings. Following Nichols & Stich (2000, pp. 124–5), I have noted that, in order for the inferential characteristics of offline imaginative states to mirror those of one's beliefs, most of one's beliefs would need to be transferred into imaginings as well. This has the odd result that we imagine almost everything we believe whenever we engage in conditional reasoning. A simpler alternative is desirable.¹³ The objector may hold that it is not really necessary that we copy all, or nearly all, of our beliefs into imagination when evaluating conditionals. Perhaps only a handful of relevant beliefs need to be tokened offline, via imaginative states, in order for the resulting inductive or abductive inferences to be suitably sensitive to our background knowledge.

While I agree that only a handful of beliefs may at any time inform a certain inference—be it online or offline—the objector faces the question of the mechanism by which only some beliefs are judged relevant to the inference at hand, so as to be copied, or not, into imagination. The problem here is that there are innumerable ways for a belief to end up being relevant (or not) to a certain inference. It is not inconceivable that there would be a sophisticated mechanism or process capable of revealing such relevance—at least in a rough-and-ready way. (Add this to the view's list of posits.) But if we are in a position to highlight for ourselves the

¹³ Nichols and Stich offer two possibilities for thinking about how propositions end up in the Possible Worlds Box (2000, pp. 124–5): on the first, all of one's beliefs are copied into the Possible Worlds Box, with the UpDater mechanism then weeding out the representations that conflict with the imagined antecedent; on the second, the UpDater acts as a “filter” only letting into the PWB representations that do not conflict with the imagined antecedent. Note that, even on the second alternative, most of one's beliefs still end up in the Possible Worlds Box, since most of one's beliefs will be consistent with most of the antecedents we in fact consider.

beliefs that are relevant to assessing a certain conditional, so as to import them into imagination, why, in addition, would we need to token the antecedent of the conditional “offline” during conditional reasoning? Why not move straight from those beliefs identified as relevant to accepting or rejecting the conditional?

6.9.2 Thought Experiments—Hard Cases for Me?

A second objection is that, at least in *some* cases, we arrive at the beliefs in the conditionals we do without relying on any background beliefs in generalizations; instead, it may be claimed, it is only by imagining a very specific scenario that we come to infer some new conditional—or, indeed, some new generalization. Such cases would challenge my claim that when we infer a new conditional of the form ‘If p , then q ,’ it is by drawing upon preexisting beliefs in generalizations and regularities relevant to p . Philosophical thought experiments would seem to be vivid examples where the details of a carefully constructed case are used to induce belief in a new conditional. In such cases, it may seem that it is precisely by imagining a very specific scenario that we first arrive at a belief in a new generalization of the sort we can then apply to future cases. Here it may seem that it is *only* by representing a certain content as true, in imagination, that one comes to believe an important new generalization—as when, for instance, we imagine a Gettier (1963) case and infer, on that basis, that justified true belief is not sufficient for knowledge. It is with this thought in mind, no doubt, that imagination—considered as a *sui generis* mental state or faculty of some kind—is invoked in discussions of the epistemic role and psychological nature of thought experiments.

In response, it is quite true that we at times come to believe a new generalization by considering a specific counterfactual scenario; and it is also true that the details of an imagined scenario are crucial in helping us to do so. But these facts pose no challenge to the idea that we only exploit existing beliefs in the process of inferring the new conditional and generalization. When imagining that p as a means to arriving at a new conditional with p as its antecedent—and whatever new generalization that may bring with it—all that matters is that we have beliefs in *some* generalizations relevant to determining what would happen if p . The value of a good thought experiment lies in its power to reveal conflicting background beliefs we didn’t know we had—and to force us to resolve them. To engage with a thought experiment is to think about what else would be true in the kind of situation described in the thought experiment, and to be pulled in opposing directions.

It will help to look at some examples. Consider, first, Searle’s (1980) famous “Chinese Room” thought experiment. It might seem to be a case where, by imagining a specific scenario, we (rightly or wrongly) come to believe a new generalization about the relation between computing and thinking. When considering Searle’s thought experiment, we ask ourselves: “If the system Searle

describes were created, would it *understand* Chinese?” Suppose that the antecedent of this question contains within it all the specifics of Searle’s thought experiment; we consider all those details when asking ourselves the question. If we react to his story in one way or the other, it will only be because considering his story brings to mind various background beliefs we have about, e.g., manipulating symbols, understanding a language, responding sensibly to prompts, and so on. Perhaps one such generalization is “The ability to respond appropriately to language prompts constitutes understanding that language”; perhaps another is “understanding requires consciousness”; perhaps a third is “non-biological systems lack consciousness.” Many philosophical thought experiments (Searle’s included) work by alerting us to the fact that two or more believed generalizations might be in conflict, by showing them to point in opposite directions with regard to a particular case. Indeed, anyone *without* beliefs in conflicting generalizations of this kind will fail to find the thought experiment interesting. The end result, for most, is revision of one of the formerly believed generalizations—or continued puzzlement. All of this is quite consistent with the claim that coming to believe “if p , then not- q ” does not require first entering into a state with the content p .

Now consider a standard Gettier case. Jones looks at a broken clock that reads 4:15 p.m. and forms the belief that it is 4:15 p.m. Coincidentally, the belief is true. We ask ourselves: what should we say about this case? Does Jones know the time? On the basis of other cases we’ve considered (pre-Gettier), we’ve come to believe that justified true belief is sufficient for knowledge. Looking at a clock seems a good way to acquire a justified belief about the time. On the other hand, we also *already* think that a belief is not knowledge when it is only coincidentally true; our having that belief is what originally led us to include justification within the set of sufficient conditions for knowledge. But now, in considering the story about Jones, it seems that our existing beliefs force us to say that his belief is true and justified, yet *also* coincidentally true. The details of the case help us to discover a dark spot in our previous belief set: we didn’t realize that the belief that knowledge is justified true belief conflicts with the belief that knowledge cannot be coincidental. Now we see the conflict and have to reconfigure. Nothing in this reasoning requires a *sui generis* state of imagining.¹⁴

¹⁴ True, within all this, we need to consume the fiction we are told about Jones (and about the Chinese room). It is commonly held that we use imagination when engaging with fictions generally. Chapters 9, 10, and 11 deal with puzzles specific to fiction. For now, I assume that simply comprehending what is going on in a fiction does not *require* imagination—that, in reading about Jones, we may simply form beliefs of the form: in the Jones fiction, thus and such. This is all that is required by the present arguments.

6.10 Recap

This chapter began with some familiar distinctions between indicative and subjunctive conditionals, using them to raise questions concerning the relevance of the Ramsey test to the psychology of how we evaluate conditionals of each sort. We saw that, when applied to subjunctive conditionals, conducting something like the Ramsey test cannot simply amount to adding a proposition hypothetically to one's belief set and revising for consistency. Accounts of imagination's role in conditional reasoning that take inspiration from the Ramsey test will require significant amendment.

I then put forward a general argument to the effect that, in any situation where the Ramsey test would take one to a true belief in a conditional, we must already have beliefs that would, by themselves, justify our inferring the conditional. So there is no normative requirement for imaginative states that would allow one to implement the Ramsey test. Next I considered whether *sui generis* imaginative states might be valuable psychological tools for arriving at beliefs in conditionals, by considering cases where imagination has been invoked to explain conditional reasoning—focusing, in particular, on Williamson's example of sheep getting out of their pen. I showed that an account involving beliefs alone is more parsimonious than Williamson's, which appeals to offline "imaginative exercises." The "belief only" account makes do without *sui generis* imaginings and requires nothing not also needed on the account involving offline imaginative exercises. Matters don't change, I argued, if we conceive of imaginative exercises as exploitations of a "skill," as opposed to inferences of a kind.

Finally, I looked at cases that might seem especially challenging for the sort of view I've advocated: philosophical thought experiments. In those cases, it might seem that we don't rely upon background generalizations when inferring a new conditional; imagining the details of the case seems to be what matters. Granting that the details matter, I argued nevertheless that, when we infer a new conditional (or generalization) in light of a thought experiment, it is because the details of the case highlight for us a conflict, or dark spot, in our preexisting beliefs. We think about what else would be true of the situation described and notice that we're pulled in opposing directions. A new belief in a conditional or generalization arises when we decide how to resolve the conflict. Understanding how all this occurs isn't aided by positing *sui generis* imaginative states.