Part II

Memory Loss

Chapter 6

Generalized Conditionalization

Up to this point we have introduced the components of a CLF model, the framework's standard interpretation, and its synchronic systematic constraints. While all of these elements have been subtly different from their traditional Bayesian forerunners, the most significant departures in CLF are in its diachronic systematic constraints, Generalized Conditionalization (GC) and the Proper Expansion Principle (PEP). These constraints were designed to help CLF model two types of certainty-loss stories: stories involving memory loss and stories involving context-sensitive claims.

While both (GC) and (PEP) are needed to model stories involving context-sensitivity, (PEP) does almost no work in modeling stories involving memory loss. So we will proceed in stages: We will first consider memory-loss stories and develop (GC); then we will consider context-sensitivity, both developing (PEP) and providing further motivation for (GC).

I begin this chapter with the Conditionalization-based framework introduced in Chapter 3—a framework that uses CLF's standard interpretation and synchronic systematic constraints but adopts Conditionalization as its diachronic constraint. This framework has served us well so far, but I will argue that when applied to stories involving memory loss or the threat thereof it yields verdicts that do not represent genuine requirements of ideal rationality. This will motivate us to replace Conditionalization and build a framework with a wider domain.

After a brief dalliance with an updating rule I call Limited Conditionalization (LC), we will settle on Generalized Conditionalization (GC) as CLF's updating constraint. I will show that (GC) properly models the memory-loss stories for which Conditionalization failed, then bolster the case for (GC) (in Section 6.2) by describing a couple of its further applications. (GC) will first be used to analyze a complex, subtle story in which an agent gains and loses information about a lottery. Then (GC) will yield an extension of van Fraassen's Reflection Principle that shows how to respond to information about a future self's degrees of belief even if that future self lacks some of your current information.

This chapter shows that (GC) yields verdicts representing substantive, interesting requirements of ideal rationality over a wide range of stories. Still, that does not explain the intuitive idea behind the updates (GC) requires. That task will be left for Chapter 7, after which we will be able to introduce (PEP) and work with CLF in its entirety.[1]

A few preliminaries: First, none of the modeling languages in this part of the book will contain sentences representing context-sensitive claims. Second, except when explicitly noted otherwise we will assume that modeling languages do not represent claims about an agent's own degrees of belief or about rational requirements on those degrees of belief. (Working with such claims would introduce all sorts of complications I wish to avoid here.) Finally, we will continue to ignore the question of how one selects the appropriate modeling language for modeling a particular story. Given a story we will adopt a model with a language that feels natural for representing it,[2] then trust that our choice will not undermine our analysis. This trust will come under scrutiny in Chapter 8.[3]

[1]Given our discussion in Section 4.2.1 of a framework's constraints as representing necessary conditions for ideal rationality, we can think of our progress over the course of this book as first building the synchronic framework, then adding (GC) to represent further necessary requirements of ideal rationality, then adding (PEP) to represent even more.

[2]Keeping in mind, as we discussed in Section 4.2.2, that the modeling language may contain sentences representing claims to which the agent does not assign a degree of belief at various times.

[3]This part of the book owes a great deal to the work of Isaac Levi, especially in his (1980). As we go along I will try to acknowledge in footnotes the points at which my arguments make the most direct contact with his, but his influence on my views has been much more extensive than can be captured by those few explicit contact points.

6.1 Updating Rules

6.1.1 Memory-Loss Objections

In Chapter 3 we introduced the following formal version of Conditionalization, the traditional Bayesian updating rule:

> **Conditionalization:** For any $t_j, t_k \in T$ with $j \leqslant k$ and any $x \in L$, $P_k(x) = P_j(x \,|\, \langle C_k - C_j \rangle)$.

Here the certainty set C_j represents the claims the agent is committed to certainty in at t_j, C_k represents the agent's committed certainties at t_k, $C_k - C_j$ is a set of representing "everything the agent learns" between those two times, and the angle-bracket notation gives us a proxy sentence for $C_k - C_j$ that is logically equivalent to the conjunction of everything in that set.

Having formalized Conditionalization, we discussed the Conditionalization-based modeling framework, a modeling framework identical to CLF except that it uses Conditionalization as its sole diachronic systematic constraint. We showed that Conditionalization-based models can achieve various intuitive results, for instance verdicts representing the requirements of ideal rationality in our story The Die.

But Conditionalization has long been known to have a peculiar property: given other standard Bayesian assumptions, if an agent updates by conditionalizing any claim she was certain of before the update will remain certain after the update. In the present context, this means that if in a Conditionalization-based model a sentence goes from an unconditional credence of 1 at an earlier time to an unconditional credence less than 1 at a later time, the model will indicate that the agent's doxastic evolution violates the requirements of ideal rationality. I prove this fact in Theorem C.3, but intuitively it holds because Conditionalization is a rule for setting degrees of belief after one has narrowed one's space of doxastic possibilities. An agent is certain of a claim at t_j just in case that claim is true in all the possibilities she entertains. When she narrows her set of doxastic possibilities between t_j and t_k, the claim will be true in all the possibilities that remain.

Talbott (1991) and Arntzenius (2003) take advantage of this fact about Conditionalization to offer memory-loss counterexamples to the updating rule. Talbott's objection to a Conditionalization-based modeling framework involves the following story:

> *Spaghetti:* At 6:30pm on March 15, 1989, Talbott is certain he is having spaghetti for dinner that night. But by March 15, 1990, Talbott has completely forgotten what he had for dinner one year ago.

A natural model of this story will have a sentence representing the claim "Talbott has spaghetti for dinner the night of March 15, 1989." There will be an extrasystematic constraint on that model assigning this claim an unconditional credence of 1 at an initial time (3/15/89), and an extrasystematic constraint assigning it an unconditional credence less than 1 at a later time (3/15/90).[4] Between those two times Talbott loses certainty in the claim represented, so a Conditionalization-based model will indicate a violation of the requirements of ideal rationality.

The lesson here is simple: If an agent suffers an episode of memory loss in which she becomes less-than-certain of a claim of which she was previously certain, a Conditionalization-based model will indicate that the agent's doxastic evolution violates the requirements of ideal rationality.[5]

Arntzenius (2003) makes the news worse for Conditionalization. He offers the following story (this version has been adapted a bit from Arntzenius's original):

> *Shangri La:* You have reached a fork in the road to Shangri La. The guardians of the tower will flip a fair coin to determine your path. If it comes up heads, you will travel the Path by the Mountains; if it comes up tails, you will travel the Path by the Sea. Once you reach Shangri La, if you have traveled the Path by the Sea the guardians will alter your memory so you remember having traveled the Path by the Mountains. If you travel the Path by the Mountains they will leave your memory intact. Either way, once in Shangri La you will remember having traveled the Path by the Mountains.
>
> The guardians explain this entire arrangement to you, you believe their words with certainty, they flip the coin, and you follow

[4]The claim that Talbott has spaghetti for dinner on 3/15/89 is neither something Talbott is certain of on 3/15/90 nor entailed by anything Talbott is certain of on 3/15/90. By the standard interpretation's Certainty Conditions (which apply to Conditionalization-based models because we assumed the Conditionalization-based framework has the same interpretation as CLF), there will be an extrasystematic constraint on our model assigning that claim a 3/15/90 credence less than 1 regardless of whether Talbott actually assigns a degree of belief to that claim on that date or not.

[5](Williamson 2000, Section 10.2) makes essentially the same point.

> your path. What does ideal rationality require of your degree of belief in heads once you reach Shangri La?[6]

As you travel to Shangri La in this story, you are certain of the outcome of the coin flip; after all, you can tell which path you are traveling. But Arntzenius argues that once you reach Shangri La, you should be uncertain whether the coin came up heads. Because of the guardians' tampering plan, you cannot rely on your memories and so cannot be certain which path you traveled to reach Shangri La.

We can analyze the Shangri La story by splitting it up into two cases: the case in which you travel the Path by the Sea and the case in which you travel the Path by the Mountains. The Path by the Sea case is much like Talbott's Spaghetti story. Between the time you are traveling the Path by the Sea and the time you are in Shangri La, memory loss causes you to lose certainty in a claim (the claim that the coin comes up tails), and a Conditionalization-based model will indicate that your doxastic evolution violates the requirements of ideal rationality. The only interesting difference between Spaghetti and the Path by the Sea case is that in the latter the memory loss results from the operations of outside agents instead of from natural processes.

Now consider the case in which the coin comes up heads and you travel the Path by the Mountains. A model of this case, model SL, is described in Table 6.1. (The extrasystematic constraints on this model incorporate Arntzenius's argument that you should be uncertain at t_2 whether the coin came up heads.) Applying Conditionalization, we can derive

$$P_2(h) = P_1(h \,|\, \langle C_2 - C_1 \rangle) \tag{6.1}$$

Yet between t_1 and t_2 no sentences go from a credence less than 1 to a credence of 1, so $C_2 - C_1$ is empty.[7] By the definition of our angle-bracket

[6] If "remember" is factive, it's not quite right to say that if you travel the Path by the Sea you will "remember" traveling the Path by the Mountains when you reach Shangri La. Perhaps the more accurate locution would be "you will have memories [or quasi-memories?] as of traveling the Path by the Mountains." I will tend to overlook this subtlety in future formulations.

[7] There are some claims in which you gain certainty between t_1 and t_2, such as "I am now in Shangri La." However, there are no claims *represented in the modeling language of SL* in which you gain certainty between t_1 and t_2, making $C_2 - C_1$ empty. One might worry that leaving the claims you learn between t_1 and t_2 unrepresented in SL will render that model's verdicts untrustworthy. In Chapter 8 we will show this is not the case by constructing a broader model of the Shangri La story whose language represents such claims; this model will yield the exact same results as model SL. That portion of the discussion is delayed until then because the extra claims are context-sensitive.

Table 6.1: Model SL

Story: Shangri La, heads case

T: Represents these times:

- t_0 After the guardians have described the process to you but before they have flipped the coin.
- t_1 While you are traveling the Path by the Mountains.
- t_2 Once you have arrived in Shangri La.

L: Built on these atomic sentences, representing these claims:

- h The coin comes up heads.
- m I travel the Path by the Mountains.

Extrasystematic constraints:

	P_0	P_1	P_2
h	<1	1	<1
$m \equiv h$	1	1	1

notation, this means that $\langle C_2 - C_1 \rangle \dashv\vdash \mathsf{T}$. So by our synchronic systematic constraints and the extrasystematic constraints on model SL, we have

$$P_2(h) = P_1(h \mid \mathsf{T}) = P_1(h) = 1 \tag{6.2}$$

But this flatly contradicts SL's extrasystematic constraint on $P_2(h)$. So this Conditionalization-based model indicates that your doxastic evolution in the Shangri La story violates the requirements of ideal rationality even if the coin comes up heads.

Yet in the heads case you suffer no actual memory loss. Since you have traveled the Path by the Mountains, the guardians leave your memory perfectly intact. It is only because of the *threat* of memory loss—the fact that when you reach Shangri La you are *uncertain* whether your memory was altered or not—that you lose your certainty at t_2 that the coin came up heads. Still, you have lost a certainty, so our Conditionalization-based model indicates that your doxastic evolution violates the requirements of ideal rationality. A negative evaluation from a Conditionalization-based model can be triggered not only by a case of actual memory loss, but also by a case in which the agent faces the *threat* of memory loss.

So far we have shown that if an agent loses certainties in a story due to memory loss or the threat thereof, a Conditionalization-based model will indicate that her doxastic evolution violates the requirements of ideal rationality. It might be suggested that this is just as it should be—a doxastic evolution that involves memory loss is less than ideal, so Conditionalization-based models are getting things right in these stories. Yet we should remember that (as we discussed in Chapter 4) ideal rationality isn't about ideality in some general sense; it concerns ideality along a particular evaluative axis, that of rational consistency. Here I think it would be highly implausible to hold that an agent's doxastic evolution can be made inconsistent just by her admitting that she *might* have forgotten something. In the Path by the Mountains Shangri La case, you go from certainty in a claim at an earlier time to less-than-certainty in that claim at a later time because you assign a positive degree of belief at that later time to the *possibility* that you have suffered memory loss. Surely we don't violate the requirements of rational consistency just by assigning a positive degree of belief to the empirical claim that we have suffered memory loss in the past, especially since such claims will often be true and well-supported by our evidence![8]

The Shangri La Path by the Mountains case demonstrates that in at least some cases the threat of memory loss will cause an agent to lose certainties without her doxastic evolution's violating the requirements of ideal rationality. Since Conditionalization-based models will indicate a violation in these stories when in fact there is none, such stories fall outside the Conditionalization-based framework's domain of applicability. Recalling the methodological points made in Chapter 5, this does not mean that Conditionalization is a failed formal updating rule. It just means that we have better delineated the sorts of stories to which models based on that rule can be fruitfully applied.

In our logical omniscience discussion (Section 5.4) we argued that rational constraints on the relations between an agent's doxastic attitudes stand or fall with rational constraints on the agent's attitudes towards logical truths. The Shangri La story ingeniously suggests that rational constraints on an agent's responses to forgetting stand or fall with rational constraints on the agent's responses to the *threat* of forgetting. If that's right, then just as an agent does not automatically violate the requirements of ideal

[8]Following on our discussion from Section 4.2.4, notice how little it helps to think about these questions in terms of ideal agents. Does such an agent's ideality endow her not only with perfect memory but also with certitude in its perfection? And even if so, what does this have to do with how rational consistency requires a real agent to assess her own memory?

rationality by suspecting she has forgotten, an agent also does not automatically violate those requirements simply by forgetting. (And after all, it isn't rationally inconsistent to *gain* information—why should it be rationally inconsistent to *lose* it?)

I suppose one could maintain that while the threat of memory loss doesn't produce a violation of the requirements of ideal rationality, actual memory loss does—so that in the Shangri La story there is a violation if you travel the Path by the Sea but not if you travel the Path by the Mountains. It's unclear what the motivation for this position would be, but someone who adopts it should still be willing to join our project of developing a new diachronic constraint that pushes beyond the boundaries of the Conditionalization-based framework's domain of applicability to include stories involving the threat of memory loss. We will soon develop a diachronic constraint that handles such stories, has intuitive applications, and (as we will see in Chapter 7) squares with appealing accounts of what it is to be doxastically consistent over time that apply even in cases in which memory loss occurs. Once we have a new diachronic constraint that models memory-loss stories and threat-of-memory-loss stories equally well, I think we can confidently say that memory loss does not violate the requirements of ideal rationality. As Williamson (2000, p. 219) puts it, "Forgetting is not irrational; it is just unfortunate."

6.1.2 Limited Conditionalization

The preceding discussion reveals that there's an important sense in which Conditionalization outstrips the intuition that motivated it. When we introduced Conditionalization in Section 3.3, we said that an agent who gains information between two times should set her later unconditional degrees of belief equal to her earlier degrees of belief conditional on what she has learned. Intuitively, that makes sense for cases in which an agent strictly gains information between two times. But Conditionalization is much broader than that: it requires an agent to update by conditionalizing between two times whether she learns information, loses information, or undergoes a combination of both.

Let's be a bit more precise. We'll refer to the net change (if any) in an agent's certanties between two times as a **doxastic event**. In a **null event**, the agent's certainties at the later time are identical to her certainties at the earlier time. In a **pure learning event**, the agent loses no certaintes; in a **pure information-loss event** the agent gains no certainties. Finally, in a **mixed doxastic event** the agent both gains and loses certainties.[9] Notice that by these definitions a null event counts as both a trivial pure learning event and a trivial pure information-loss event. Notice also that these definitions work in terms of the *net* change in the agent's certainties; if for instance an agent loses all her certainties and then gains them back again between two times, the doxastic event that runs from one of those times to the other counts as a null event.

When we motivated Conditionalization intuitively, we did so by thinking about pure learning events, such as the one that happens to Marilynn in The Die. But when we try to apply the Conditionalization-based framework to stories like Shangri La in which agents lose certainties, its models yield verdicts that do not represent requirements of ideal rationality. The domain of applicability of the Conditionalization-based framework is limited to the sorts of stories that originally motivated it: stories in which all the doxastic events are pure learning events.[10]

We want a diachronic constraint that will give our modeling framework a broader domain—that will allow it to successfully model stories in which agents both gain and lose certainties. But in expanding our framework's purview we don't want to lose the fruitful verdicts that conditionalizing yields. Since conditionalizing makes sense when an agent undergoes a pure learning event, our diachronic constraint should allow conditionalizing in those cases. Here's a constraint that does that:

> **Limited Conditionalization (LC)**: For any $t_j, t_k \in T$ with $j \leqslant k$ and any $x \in L$, if $C_j \subseteq C_k$ then $P_k(x) = P_j(x \mid \langle C_k - C_j \rangle)$.

Limited Conditionalization looks exactly like Conditionalization, except that it applies only when $C_j \subseteq C_k$. That is, (LC) requires an agent's later

[9] Technically the only way to class a doxastic event as one of these kinds is to choose a model that represents the two times and check whether any of the claims represented in that model's language go from certainty to less-than-certainty (or *vice versa*) between those times. So a description of a doxastic event as a pure learning event, say, is always relative to a choice of modeling language. But since we have set aside questions of language choice until Chapter 8, until then we will categorize doxastic events against an implicit relativization to whatever natural-seeming language we have adopted to represent the story being analyzed.

[10] It's very possible that this domain restriction was already implicit in the practice of decision theorists and statisticians. Schervish, Seidenfeld, and Kadane (2004) respond to Arntzenius that it "is already assumed as familiar in problems of stochastic prediction" that conditionalization updating rules are to be applied only when the agent's certainty sets form a filtration—that is, when no certainties are lost at any point during a story.

degrees of belief to be a conditionalization of her earlier degrees of belief only when she retains all certainties at the later time that she had at the earlier time. In other words, (LC) authorizes conditionalizing only for pure learning events.[11]

Suppose we build a modeling framework that has CLF's standard interpretation and synchronic systematic constaints, but that takes (LC) as its diachronic systematic constraint. In other words, suppose we start with the Conditionalization-based framework and then substitute (LC) for Conditionalization. Does this substitution allow us to model Shangri La?

The first improvement here is that an (LC) analysis of model SL will not yield Equation (6.2), which indicated that you should be certain once you reach Shangri La that the coin came up heads. Equation (6.2) derived your t_2 degrees of belief from your t_1 degrees of belief by Conditionalization. But since you lose certainties between t_1 and t_2 (such as your certainty that you traveled the Path by the Mountains), $C_1 \nsubseteq C_2$. So (LC) does not require you to generate your t_2 degrees of belief from your t_1 degrees of belief by conditionalizing—in fact, (LC) doesn't directly relate t_1 and t_2 credences at all! If we use (LC) as our diachronic constraint, model SL will not have the same problem it had under a Conditionalization-based regime; it will not indicate that your doxastic evolution violates the requirements of ideal rationality. So the Shangri La story falls within an (LC)-based framework's domain of applicability.

The move from Conditionalization to (LC) keeps us from having to conditionalize in doxastic events that are not pure learning events. In a sense, it's a defensive move, preventing our models from making certain mistakes—from yielding incorrect verdicts relating credences at two times. But surprisingly, (LC) can also go on offense; it can reveal positive requirements of ideal rationality that we couldn't see in our models before.

To see how, consider the relation in Shangri La not between t_1 and t_2 but between t_0 and t_2. Even though you gain and lose certainties as your path unfurls between those two times, the net effect of those certainty changes cancels out. Your certainty set at t_2 is identical to your certainty set at t_0, so the doxastic event you undergo between those two times is a null event—the trivial case of a pure learning event. Since $C_0 \subseteq C_2$, (LC) will relate your t_2 certainties to your t_0 certainties by conditionalizing. And since $C_2 - C_0$

[11] Corollary C.2 guarantees that if $C_j \subseteq C_k$ then $P_j(\sim\langle C_k - C_j\rangle) < 1$, so we need not worry than (LC) will set an unconditional credence equal to an undefined conditional credence.

is empty, an (LC)-based version of SL will yield:

$$P_2(h) = P_0(h \,|\, \mathsf{T}) = P_0(h) \tag{6.3}$$

This verdict indicates a genuine requirement of ideal rationality. The Shangri La story asks what your degree of belief that the coin came up heads should be once you reach Shangri La. While you were on whichever path you traveled, you knew for certain the outcome of the coin flip. But once you reach Shangri La, your evidential situation with respect to the coin (as represented in your certainty set) is reset to precisely what it was before you began your trip. Thus your degree of belief that the coin came up heads should revert to what it was before the coin was flipped. If, for instance, the Principal Principle is a requirement of ideal rationality, you should assign a 1/2 degree of belief to the claim that the coin comes up heads both at t_0 and at t_2.[12]

(LC) also reveals an interesting fact about your t_2 degrees of belief in Shangri La. Since C_0 and C_2 are identical, an (LC)-based model will yield a verdict of the following form for any $x \in L$:

$$P_0(x) = P_2(x) \tag{6.4}$$

Since $C_0 \subseteq C_1$, for any $y \in L$ (LC) will yield:

$$P_1(y) = P_0(y \,|\, \langle C_1 - C_0 \rangle) \tag{6.5}$$

By Lemma C.1, $P_0(\sim\langle C_1 - C_0 \rangle) < 1$, so by the Ratio Formula

$$P_1(y) = \frac{P_0(y \,\&\, \langle C_1 - C_0 \rangle)}{P_0(\langle C_1 - C_0 \rangle)} \tag{6.6}$$

Applying Equation (6.4) twice,

$$P_1(y) = \frac{P_2(y \,\&\, \langle C_1 - C_0 \rangle)}{P_2(\langle C_1 - C_0 \rangle)} \tag{6.7}$$

The Shangri La story is set up so that between t_1 and t_2 you lose precisely those certainties you gained between t_0 and t_1. Thus we have $\langle C_1 - C_0 \rangle \dashv\vdash$

[12] A similar point could be made about the Spaghetti story. Suppose Talbott's dinner entree is selected for him during the day by a chance process whose probability of picking spaghetti on a given night does not change from one year to the next; Talbott is certain throughout the story what this objective probability is and that it does not change. In that case, ideal rationality requires him to assign the same March 15, 1990 degree of belief to the claim that he had spaghetti on March 15, 1989 as he assigned to that claim when he woke up on March 15, 1989. This requirement would be indicated by a verdict of an (LC)-based model of Spaghetti.

$\langle C_1 - C_2 \rangle$. Applying our Equal Unconditional Credences in Equivalents principle twice, we have

$$P_1(y) = \frac{P_2(y \,\&\, \langle C_1 - C_2 \rangle)}{P_2(\langle C_1 - C_2 \rangle)} \tag{6.8}$$

Finally, since $P_2(\langle C_1 - C_2 \rangle) = P_0(\langle C_1 - C_0 \rangle)$ and the latter is positive, we can apply the Ratio Formula to obtain

$$P_1(x) = P_2(x \,|\, \langle C_1 - C_2 \rangle) \tag{6.9}$$

Take a moment to compare Equations (6.5) and (6.9). According to an (LC)-based model, your degrees of belief at t_2 are required to have the same relation to your degrees of belief at t_1 as your degrees of belief at t_0 had to your degrees of belief at t_1.[13]

The modeling frameworks we are examining model changes in an agent's degrees of belief driven by changes in the claims she takes for certain. Between t_0 and t_1 you experience a pure learning event and respond by updating your degrees of belief by conditionalization. That is, your unconditional degrees of belief at the later time equal your degrees of belief at the earlier time conditional on the certainties gained in-between. Between t_1 and t_2 you experience a pure certainty-loss event in which you lose all the certainties you gained between t_0 and t_1. From the point of view of your certainty set, the latter event is just the former happening backwards in time. So it should come as no surprise that your t_1 degrees of belief are related to your t_2 degrees of belief by a *reverse-temporal* conditionalization: your unconditional degrees of belief at the *earlier* time equal your degrees of belief at the *later* time conditional on the certainties *lost* in-between.[14]

6.1.3 (GC)

A rule requiring reverse-temporal conditionalization in response to pure certainty-loss events could be supported by the following principle: when you lose information, your resulting doxastic state should be such that were you to regain that information you would return to the doxastic state in which you began. But if we hold that principle in general, we are going to have change updating rules again.

[13](Levi 1987, p. 198) presents an argument of precisely this form for the general case of a pure certainty-loss event. (I am grateful to Teddy Seidenfeld for bringing this argument to my attention.)

[14]cf. Levi's discussion of "inverse temporal credal conditionalization" at (Levi 1980, Section 4.3).

We obtained Equation (6.9) from (LC) indirectly, by first relating t_2 credences to t_0 credences and then relating t_0 credences to t_1 credences. Our strategy for relating P_1 to P_2 relied on there being a time in the story before t_1 at which the agent had a certainty set identical to C_2. But imagine we have a story with only two times, between which a pure certainty-loss event occurs. There will be no t_j and t_k ($j < k$) available such that $C_j \subseteq C_k$, so the (LC)-based framework will be unable to yield any diachronic verdicts for this story. Still, we may think that ideal rationality requires the agent's doxastic state after the pure information-loss event to be such that if she re-learned the information lost and conditionalized upon it she would return to her initial state.

To obtain a verdict from our modeling framework representing this requirement of ideal rationality, we need a stronger updating rule than (LC). CLF adopts

> **Generalized Conditionalization (GC)**: For any $t_j, t_k \in T$ and any $x \in L$, if $P_j(\sim\langle C_k - C_j\rangle) < 1$ and $P_k(\sim\langle C_j - C_k\rangle) < 1$ then
> $$P_j(x \,|\, \langle C_k - C_j\rangle) = P_k(x \,|\, \langle C_j - C_k\rangle)$$

To get a rough intuitive idea of what (GC) says, let's refer to the t_j agent as "Jill" and the t_k agent as "Ken". (GC) says that if Jill supposes everything Ken is certain of but she isn't, and Ken supposes everything Jill is certain of but he isn't, the two of them are required to assign the same degree of belief to any claim. The two antecedent conditions in (GC) are there to guarantee that the relevant conditional credences are defined.[15]

Theorem C.5 demonstrates that (GC) is as strong as (LC); every verdict yielded by an (LC)-based model (such as our (LC) verdicts for Shangri La) will be obtainable from a (GC)-based model as well. But Theorem C.5 also demonstrates that (GC) entails this reverse-temporal twin of (LC):

> For any $t_j, t_k \in T$ with $j \leqslant k$ and any $x \in L$, if $C_k \subseteq C_j$ then $P_j(x) = P_k(x \,|\, \langle C_j - C_k\rangle)$.

[15] These antecedent conditions are framed in terms of sentences' having unconditional credences less than 1 both so they will mesh with the antecedent in CLF's version of the Ratio Formula and so that (GC) can be used in derivations. To introduce an instance of (GC) as a line in a derivation, we first determine which sentences are in $C_k - C_j$ and $C_j - C_k$ and whether $P_j(\sim\langle C_k - C_j\rangle)$ and $P_k(\sim\langle C_j - C_k\rangle)$ are both less than 1. These facts can be determined from extrasystematic constraints that will already be listed as premises in our derivation.

So a (GC)-based framework (that is, a framework with CLF's standard interpretation and synchronic systematic constraints but with (GC) as its diachronic constraint) will give us reverse-temporal conditionalization relations even in stories with a single, pure certainty-loss doxastic event. A (GC)-based framework is strictly stronger than an (LC)-based framework, in the sense that it yields all of the latter's verdicts while also yielding more besides. Yet (GC) does not yield the bad sorts of verdicts we got from Conditionalization in the Shangri La story; (GC) does not, for instance, yield Equation (6.2) indicating that you are required once you reach Shangri La to be certain that the coin came up heads. Shangri La (and Spaghetti, for that matter) lies within a (GC)-based framework's domain of applicability.

I view (GC) as a *generalization* of traditional conditionalization constraints because it relates credences separated by pure learning events (yielding forward-temporal conditionalizations), by pure certainty-loss events (yielding reverse-temporal conditionalizations), and by mixed doxastic events (yielding relations between conditional credences at the two times). Of course, one should not simply take my word that (GC) is a useful generalization. The rest of this chapter will describe stories for which (GC) yields verdicts representing genuine requirements of ideal rationality. In the next chapter, we will describe a deeper intuition underlying (GC), which I think provides strong explanations of why conditionalizing seemed like a good idea for pure learning events in the first place. Before moving on, however, I'd like to present some alternative, equivalent forms of (GC).[16]

According to Theorem C.8, we can obtain a diachronic constraint equivalent to (GC) by replacing its consequent with either of the following:

$$P_j(x \,|\, \langle C_k \rangle) = P_k(x \,|\, \langle C_j \rangle) \tag{6.10}$$

or

$$P_j(x \,|\, \langle C_j \cup C_k \rangle) = P_k(x \,|\, \langle C_j \cup C_k \rangle) \tag{6.11}$$

Moreover, Theorem C.9 demonstrates that in a consistent model the antecedent of (GC) is met just in case $P_j(\sim\langle C_j \cup C_k \rangle) < 1$ and $P_k(\sim\langle C_j \cup C_k \rangle) < 1$, which in turn is true just in case $C_j \cup C_k$ is consistent.[17] By swapping in and out various antecedents and consequents, we can obtain equivalent forms of (GC). I have chosen the one originally presented as the "official" version because of its formal continuity with updating rules

[16] I am grateful to Carl Wagner for suggesting that I include alternative (and perhaps more intuitive) formulations of (GC).

[17] If a model is inconsistent it's irrelevant what form the antecedent of (GC) takes, because the consequent will be a verdict of the model no matter what.

like Conditionalization and (LC). Also, the official version tends to be the quickest to work with when deriving verdicts, because one need only list the *differences* between C_j and C_k instead of listing the contents of each certainty set or the contents of their union. On the other hand, when proving theorems it is often easiest to work with the following (GC) equivalent:

> **Formal (GC):** For any $t_j, t_k \in T$ and any $x \in L$, if $C_j \cup C_k$ is consistent then $P_j(x \,|\, \langle C_j \cup C_k \rangle) = P_k(x \,|\, \langle C_j \cup C_k \rangle)$.

Formal (GC) will be used in various technical arguments and proofs to come. It also reminds us of an important fact that we will use repeatedly later on: (GC) yields verdicts relating credences at two times just in case the certainty sets indexed to those times are logically consistent with each other.

6.2 Applications of (GC)

6.2.1 The Lottery

Shangri La is a fairly simple story. You experience a pure learning event, then you experience a pure certainty-loss event that sets your certainties exactly back to their original state. No wonder your non-extreme degrees of belief at the end of the story are required to match their values at its beginning. But how does the (GC)-based framework fare in subtler situations involving mixed doxastic events and agents who aren't returned to their initial certainty states?

To get an idea, consider the following story:

> *The Lottery:* A lottery is being held for an *enormous* cash prize. Ten final contestants have been chosen, and each name has been placed into a hat. To heighten the drama, contestant names will be drawn out one by one (in such a manner that at each stage each remaining name has an equal chance of being picked next) and announced by a different celebrity each time. The last contestant whose name remains in the hat will win the prize.
>
> Dave is one of the ten contestants in this lottery. Going into the contest, he assigns a 1/10 degree of belief to each of the contestant's prospects. The drawing gets going, with Hugh Jackman announcing Kim's name first. Gradually more names are called, until only Al, Dave, and Frank remain. At this moment, Dave excitedly assigns a 1/3 degree of belief to each of his, Al's, and Frank's prospects. Then Nicole Kidman announces Frank's

> name. Dave is elated, and assigns a 1/2 degree of belief that he will be the eventual winner.
>
> Dave gets to thinking about what he will do if he wins all that money, and his mind wanders. Some time later, he finds he has forgotten whose name Ms. Kidman announced. He is certain of the seven names that were eliminated first, and he is certain that the last name announced wasn't his, but he now assigns a 1/2 degree of belief that it was Al's name and a 1/2 degree of belief that it was Frank's. Still, Dave retains his 1/2 degree of belief that he will be the winner, and distributes the remaining 1/2 equally between Al's and Frank's prospects of victory.
>
> The lottery organizers fuss around a while longer, and Dave's mind wanders a bit more. Eventually he snaps his attention back to the present, and to remind himself where things stand he mentally lists off all the contestants who have been eliminated. In doing so, he forgets to include Kim's name on the list. So while Dave remains certain he hasn't been eliminated, he now has six names that he was sure were announced plus the recognition that he can't remember whose name Nicole Kidman just announced. Dave thinks there are three names left in the hat (including his own), and assigns a 1/3 degree of belief to Al's, Frank's, or Kim's names being the one most recently announced. Nevertheless, Dave maintains his 1/2 degree of belief that he will be the winner, distributing the remaining 1/2 equally among the prospects of the other contestants he thinks may still be in the running.

The Lottery is a complicated story, but I think it nicely illustrates some subtle features of the (GC)-based framework. Before we discuss its (GC) analysis, take a few moments to work through the story and decide whether Dave's doxastic evolution violates the requirements of ideal rationality, and if so at what point it begins to go off the rails. To make matters easier, Table 6.2 displays a partial (GC)-based model of The Lottery that lists a time set and extrasystematic constraints reflecting the degrees of belief assigned by Dave at various times in the story.[18] The rest of this model is long and intricate, so I'll simply present its results here.

The first interesting thing to note is that none of Dave's individual doxastic states violates the requirements of ideal rationality represented in CLF's

[18] In Table 6.2 the "Kidman announces. . . " claims are meant to be tenseless, "eternal" claims of the form "Nicole Kidman has announced, is announcing, or will announce. . . ." This prevents those claims from being context-sensitive.

Table 6.2: A Partial Model for The Lottery

T: Represents these times:

- t_0 Before any of the ten names is announced.
- t_1 After seven names have been announced; only Al, Dave, and Frank remain.
- t_2 Just after Frank's name is announced (only Al and Dave remain).
- t_3 After Dave has forgotten that Al's name was the last one announced. Dave is certain that only two names remain (including his own), but can't remember whether Al's or Frank's is the other name in the hat.
- t_4 After Dave has created his mistaken list; he is certain there are three names in the hat (including his), but can't remember which two of Al's, Frank's, and Kim's join his.

Extrasystematic constraints:

	P_0	P_1	P_2	P_3	P_4
Dave will win.	1/10	1/3	1/2	1/2	1/2
Al will win.	1/10	1/3	1/2	1/4	1/6
Frank will win.	1/10	1/3	0	1/4	1/6
Kim will win.	1/10	0	0	0	1/6
Kidman announces Dave's name.	1/10	1/3	0	0	0
Kidman announces Al's name.	1/10	1/3	0	1/2	1/3
Kidman announces Frank's name.	1/10	1/3	1	1/2	1/3
Kidman announces Kim's name.	1/10	0	0	0	1/3

synchronic constraints. Each credence function $P_i(\cdot)$ represented as a column in Table 6.2 satisfies all of Kolmogorov's probability axioms.[19] Second, if a (GC)-based model of The Lottery includes only times t_0 through t_3 in its time set, it does not indicate any violation of the requirements of ideal rationality. But when t_4 is added to the time set, a (GC)-based model does indicate a violation of the requirements of ideal rationality. The (GC)-based framework indicates that if Dave assigns the t_4 degrees of belief described in Table 6.2 after assigning the other degrees of belief described there, his doxastic evolution violates the requirements of ideal rationality.

And this is exactly as it should be. Dave's transition from t_2 to t_3 is perfectly rational; while he can no longer remember at t_3 whose name was just eliminated, he remains certain that there are only two names left in the hat, so it is rational for him to maintain his 1/2 degree of belief that he'll be the winner. But between t_3 and t_4, when Dave goes from thinking there are two contestants left in the lottery to thinking there are three, that should decrease his confidence in his own prospects.

The Conditionalization-based framework would not be able to reproduce these results. That framework indicates a violation the moment an agent loses any certainties, so a Conditionalization-based model would indicate that Dave's doxastic evolution violated the requirements of ideal rationality even if that evolution included only t_0 through t_3.

The (GC)-based framework, on the other hand, gets The Lottery right. (As does (LC), for that matter.) This is despite the fact that (as one can tell by inspecting the distribution of zeroes in Table 6.2) Dave's certainty sets at t_3 and t_4 are different from the certainty sets he possessed at any other time during the story. Also, if we look at just t_1 and t_4, the overall transition from earlier time to later time is a mixed doxastic event. (Dave has certainties at t_4 he didn't have at t_1, and *vice versa*.) Even in a complex story with mixed doxastic events and certainty-loss events that don't return an agent to a previous condition, the (GC)-based framework correctly models requirements relating an agent's evolving doxastic states.

[19] If you are skeptical that the t_4 degrees of belief described for Dave can be synchronically consistent, carry out the following process for filling out Dave's t_4 degrees of belief in claims of the form "X will win and Kidman announces Y's name": First, go through and assign degrees of belief of 0 to the claims of this form Dave rules out at t_4. Then assign degree of belief 1/12 to each of the claims that remain except for claims that begin "Dave will win and...." Finally, assign each of the "Dave will win and ..." claims a degree of belief of 1/6. The result is a degree of belief distribution representable by a probabilistic credence function that yields all the values in the last column of extrasystematic constraints in Table 6.2.

6.2.2 (GC) and Reflection

The initial target of Talbott's and Arntzenius's memory-loss objections is actually not an updating rule; it is van Fraassen's Reflection Principle (1995). Conditionalization rules get dragged into Talbott's and Arntzenius's discussions largely because they can be used to justify the Reflection Principle. In this section I want to examine Talbott's and Arntzenius's objections to Reflection and suggest a (GC)-based Generalized Reflection Principle that avoids them.

To do this, however, we have to step somewhat off the reservation methodologically. Arguments for the Reflection Principle involve second-order degrees of belief and the agent's reasoning explicitly about what ideal rationality requires of her. For the reasons noted in Section 5.3.3, claims entertained by an agent about her own degrees of belief and about her own rationality may create special problems when represented in CLF. So this section will present informal arguments that ideal rationality requires particular degrees of belief of agents, instead of deriving those requirements formally in CLF. These arguments involve some very subtle conditions concerning an agent's beliefs about her own beliefs, and their informal nature makes them difficult to check. So while I endorse the results, I can't completely guarantee them.

Since our arguments will be conducted informally and not within a CLF model, they will be arguments directly about relations between an agent's degrees of belief in claims (instead of being formal derivations about relations between credences in sentences). To simplify presentation I will borrow CLF's notation, but employ it differently than in the rest of this book. For the rest of this section strings of italicized letters will refer to claims—they will be *names* for claims, instead of being sentences in a model representing those claims. An equation involving the expression $P_i(x)$ will describe a property of the agent's degree of belief in claim x at t_i. C_i will refer to the set of claims the agent is committed to certainty in at t_i (that is, if $x \in C_i$ then the agent violates the requirements of ideal rationality if she assigns less-than-certainty to x at t_i). $\langle S \rangle$ will be a claim equivalent in sentential logic to the conjunction of the claims in set S. After this section our notation will revert to its standard usage.

Using this notation, we can express the Reflection Principle as follows:

> **Reflection**: Suppose we have a claim x and two times t_j and t_k with $j \leqslant k$. For some real $\mathbf{r}$, let f be the claim that $P_k(x) = \mathbf{r}$. Ideal rationality requires $P_j(x \mid f) = \mathbf{r}$.

The rough idea of the Reflection Principle is that to the extent she can an agent should at an earlier time defer to her unconditional degrees of belief at a later time. For example, it follows from Reflection (and other synchronic requirements of ideal rationality) that if the agent is certain at t_j that she will assign unconditional degree of belief $\mathbf{r}$ to x at future time t_k, she should assign $\mathbf{r}$ to x at t_j as well. In line with our discussion in Section 4.2.2, we should read the last sentence of Reflection as stating that if the agent assigns $P_j(x \mid f) \neq \mathbf{r}$, her doxastic state at t_j violates the requirements of ideal rationality. (No violation occurs if the agent fails to assign a $P_j(x \mid f)$ value at all.)

What happens when we apply Reflection to Shangri La? Let f be the claim that at t_2 (once you have arrived in Shangri La) you assign a degree of belief of 1/2 to the claim that the coin came up heads (which we'll call h). At t_1 you are certain of f, because you know that once you reach Shangri La the threat of memory tampering will leave you uncertain which path you traveled to get there.[20] So according to the Reflection Principle, ideal rationality requires $P_1(h) = 1/2$. But whichever path you are on at t_1, you are certain at that time that you are on that path, so you should either be certain that the coin came up heads or certain that the coin came up tails. At t_1 ideal rationality requires you to assign an extreme degree of belief to h—Reflection gets the Shangri La story wrong.

This result is unsurprising for two reasons. It is generally acknowledged that given a few side conditions (similar to ones we'll presently enumerate) Conditionalization entails Reflection. Van Fraassen also argues that Reflection entails Conditionalization.[21] These tight connections make it unsurprising that Reflection renders incorrect judgments for memory loss stories just as Conditionalization does. Second, one good intuitive reason for an agent to defer to her future self's degrees of belief in the manner suggested by Reflection is that her future self typically has more information than she does at present.[22] But in stories like Shangri La, in which the agent is certain that her future self will lack information she currently possesses, ideal rationality does not require her to defer to that future self's degrees of belief.[23]

[20] We'll assume for simplicity's sake that at earlier times during the Shangri La story you are certain that you actually will assign degrees of belief to a variety of claims at later times.

[21] (van Fraassen 1995); though see also (Weisberg 2007). Generally, the technical arguments in this section have greatly benefitted from the discussions in Weisberg and van Fraassen.

[22] See, for instance, (Evnine 2008, Chapter 5) for discussion of this thought.

[23] This point is hardly original to Talbott or Arntzenius. For example, in his (1987, p. 204), Levi notes that Reflection rules out "contraction" cases in which an agent loses

With its close ties to Conditionalization, the Reflection Principle works well for stories in which all the doxastic events are pure learning events. To obtain a more general version of Reflection suitable for all types of doxastic events, we need a principle based on (GC). The key is to focus not on the agent's current suppositions about her future *unconditional* degrees of belief, but instead on her suppositions about her future self's *conditional* degrees of belief. Our more general principle runs as follows:

> **Generalized Reflection**: Suppose we have a claim x and two times t_j and t_k. For some real $\mathbf{r}$, let f be the claim that $P_k(x \mid \langle C_j - C_k \rangle) = \mathbf{r}$. Ideal rationality requires $P_j(x \mid f) = \mathbf{r}$.

Again, we should read the last sentence of this principle as stating that if the agent assigns $P_j(x \mid f) \neq \mathbf{r}$, her doxastic state at t_j violates the requirements of ideal rationality. Theorem C.14 proves that Generalized Reflection entails Reflection in cases in which the agent is certain at t_j that $C_j \subseteq C_k$ and that her t_k degrees of belief will satisfy the requirements of ideal rationality represented in CLF's synchronic systematic constraints.

The intuitive idea of Generalized Reflection is this: If the agent believes at t_j that her t_k self will lack some information (certainties) that she possesses at t_j, the t_j agent should not defer to the t_k agent's unconditional degrees of belief. But if the t_k agent were to gain back all the certainties lost since t_j (that is, $C_j - C_k$), she would then be at least as well informed as the t_j agent and worth deferring to. Since ideal rationality requires agents to update by conditionalizing in response to pure learning events, the degrees of belief the t_k agent would assign if she learned the claims in $C_j - C_k$ equal her t_k degrees of belief conditional on the supposition of $\langle C_j - C_k \rangle$. So at t_j the agent should defer not to the t_k agent's unconditional degrees of belief, but instead to the t_k agent's degrees of belief conditional on $\langle C_j - C_k \rangle$.[24]

If we assume that (GC) (as interpreted under CLF's standard interpretation) represents a requirement of ideal rationality, Generalized Reflection can be derived from it under particular conditions. Theorem C.12 shows that given a claim x, two times t_j and t_k, and a real number $\mathbf{r}$, if f is the claim $P_k(x \mid \langle C_j - C_k \rangle) = \mathbf{r}$ then Generalized Reflection will hold when the following conditions are met:

---

certainties between t_j and t_k.

[24]Generalized Reflection is very close to a principle proposed in (Elga 2007); the paragraph to which this note is attached provides roughly Elga's explanation of that principle. On the other hand, Generalized Reflection bears almost no relation to the unfortunately-similarly-named "General Reflection" principle in (Weisberg 2007).

1. At t_j the agent is certain of the claim "The doxastic evolution consisting of my t_j and t_k doxastic states satisfies all the requirements of ideal rationality."

2. At t_j the agent is certain of the claim "All the claims in C_j are true, and all the claims in C_k are true."

3. $P_j(f) > 0$.

4. At t_j the agent can identify a finite set of claims E such that:

    (a) At t_j the agent is certain of the claim "For any distinct $y, z \in E$, $\sim(y \,\&\, z)$."

    (b) At t_j the agent is certain of the claim "$\langle C_k - C_j \rangle$ is in E."

    (c) For each $y \in E$, the agent assigns a degree of belief at t_j to x conditional on y and is certain what that degree of belief is.

These conditions are sufficient to prove Generalized Reflection from (GC); they are not necessary.[25] However, something in the vicinity of Conditions 1 and 2 must hold in a story for Generalized Reflection to express a requirement of ideal rationality in that story. Condition 1 is important because the t_j agent would not want to automatically defer to her t_k degrees of belief if she suspected those degrees of belief were assigned irrationally. As for Condition 2, if the agent suspects some of the certainties she gains between t_j and t_k will be false, she will not want to automatically defer to t_k degrees of belief based upon those certainties.[26] Conditions 1 and 2 combine to rule out examples in which the t_j agent believes she will be drunk at t_k, or will be under the effects of Christensen's (1991) hypothetical psychedlic drug LSQ.[27] (Condition 3 covers the case in which, given Conditions 1 and 2,

[25]For example, one can weaken Condition 4b to the agent's being certain at t_j that $\langle C_k - C_j \rangle$ will turn out to be logically equivalent to some claim in E. One can also weaken Condition 4c to the agent's being *committed* at t_j to a particular degree of belief in x conditional on y and being certain of what that committed value is. Further weakenings may be possible, but I find that they tend to ratchet up the complexity of the proof without making the result much more informative.

[26]In Section 9.1.3 I will explain how exactly Condition 2 should be read when we are dealing with context-sensitive claims.

[27]Depending on one's theory of what ideal rationality requires an agent to believe about her own current degrees of belief and certainties, it may be redundant to require in Condition 1 that the agent be certain at t_j that her t_j degrees of belief are rational, or to require in Condition 2 that the agent be certain at t_j that all the claims in C_j are true. Since these points are required to make the proof go through, I have chosen to make them explicit in Conditions 1 and 2.

the agent is certain that her t_k self will not assign $P_k(x \mid \langle C_j - C_k \rangle) = \mathbf{r}$, so $P_j(x \mid f)$ is undefined.)

I have put claims in quotes in the conditions to clarify various use/mention issues. For example, at t_j the agent is not certain which claims will be in C_k. Condition 2 does not say of whatever claims are in C_k that the agent is certain at t_j that those claims are true. Instead, the agent is certain at t_j that whatever claims turn out to be in C_k, those claims will be true. A similar point goes for the claim f referred to in Generalized Reflection and Condition 3. The agent may not be certain at t_j what claims she will lose certainty in between t_j and t_k. Thus claim f does not say of whatever claims are in $C_j - C_k$ that the agent's t_k degree of belief in x conditional on those claims is $\mathbf{r}$. Instead, claim f says "At t_k I will assign to x degree of belief $\mathbf{r}$ conditional on a claim logically equivalent to the conjunction of the certainties I lose between now and then."[28]

Condition 4 states that the agent can construct a partition representing the different possible information sets she may gain between t_j and t_k. The partition will be a set of claims such that the agent is certain that they are mutually exclusive and is certain that one of them will be a proxy for $C_k - C_j$. This might happen, for example, if the agent is going to observe an experiment between t_j and t_k and can at t_j enumerate all the distinct possible experimental outcomes.

The rough idea of the proof is that when the agent supposes claim f, she supposes that $P_k(x \mid \langle C_j - C_k \rangle) = \mathbf{r}$. Since she is certain that her P_k values will be related to her P_j values by (GC) (as stated in Condition 1), this is to suppose that $P_j(x \mid \langle C_k - C_j \rangle) = \mathbf{r}$. The agent can then comb through the claims in the partition and collect those conditional on which she currently assigns degree of belief $\mathbf{r}$ to x. (Condition 4c ensures she has the information she needs to do this.) To suppose that $P_j(x \mid \langle C_k - C_j \rangle) = \mathbf{r}$ is to suppose that $\langle C_k - C_j \rangle$ is one of the claims in this collection. Since the agent is certain that the claims she learns between t_j and t_k are true (Condition 2), supposing that one of the claims in her collection is equivalent to $\langle C_k - C_j \rangle$ is tantamount to supposing that one of those claims is true. And conditional on that supposition ideal rationality requires her to assign degree of belief $\mathbf{r}$ to x.

Reflection suffered when we applied it to the Shangri La story—in par-

[28] The agent will also be uncertain at t_k which claims are in $C_j - C_k$. When she assigns the fateful conditional degree of belief—$P_k(x \mid \langle C_j - C_k \rangle)$—at t_k she does not assign it under that description. Instead, among all the other degrees of belief she assigns at t_k she innocently assigns a conditional degree of belief $P_k(x \mid y)$ for some y which turns out, unbeknownst to her, to be logically equivalent to the conjunction of the claims in $C_j - C_k$.

ticular when we applied it to your degrees of belief at t_1 (while you are traveling to Shangri La). How does Generalized Reflection fare? Shangri La lies in the domain of applicability of (GC), and Generalized Reflection can be derived from (GC) under particular conditions, so Generalized Reflection ought to provide correct judgments about the requirements of ideal rationality in Shangri La as long as the conditions are met.

For the sake of definiteness, let's focus on the case in which you travel the Path by the Mountains (the analysis would run symmetrically for the Path by the Sea case). To apply Generalized Reflection, we need to focus on a claim f that concerns your t_2 degrees of belief conditional on the claims you lose certainty in between t_1 and t_2. The relevant claim f takes the form $P_2(h \,|\, \langle C_1 - C_2 \rangle) = \mathbf{r}$. Since $\langle C_1 - C_2 \rangle \dashv\vdash h \,\&\, m$, suppose we assign $\mathbf{r} = 1$, letting f be $P_2(h \,|\, h \,\&\, m) = 1$. We can assume that Conditions 1 and 2 are met at t_1 with respect to t_2. Condition 3 is met, since by your t_1 assumption that you meet the requirements of ideal rationality at t_2 (Condition 1), you are certain of f. Condition 4 is met at t_1 by the set E that contains only a tautology (because $C_2 - C_1$ is empty). Generalized Reflection therefore says that ideal rationality requires you to assign $P_1(h \,|\, f) = 1$, which is obviously correct since you assign t_1 degrees of belief of 1 to h, f, and $h \,\&\, f$.[29]

We have been talking about Generalized Reflection as if t_k is always a later time than t_j. But like (GC), Generalized Reflection does not require $j \leqslant k$. And this makes sense: if we are working with stories in which an agent can lose certainties, the present agent may want to defer to an earlier version of herself who possessed information she currently lacks. For example, in Shangri La we can apply Generalized Reflection to your t_2 degrees of belief concerning the degrees of belief you assigned at t_1. Let t_j be t_2, t_k be t_1, and f be $P_1(h \,|\, \langle C_2 - C_1 \rangle) = 1$. At t_2 you are certain that $\langle C_2 - C_1 \rangle \dashv\vdash \mathsf{T}$, so you are certain that f is true just in case $P_1(h) = 1$, which in turn is true just in case h is. By Substitution (see Section 3.2.5) ideal rationality requires $P_2(h \,|\, f) = P_2(h \,|\, h) = 1$, just as Generalized Reflection predicts.

The Reflection principle tells an agent how to respond to suppositions about the degrees of belief of future versions of herself who have at least as much information as she does. We now have a Generalized Reflection principle that tells an agent how to respond to suppositions about versions of herself who may both possess information she currently lacks and lack information she currently possesses. Moreover, those versions of herself may be either in the past or in the future. This is a much more general principle, but under the right conditions it is provable directly from (GC).

[29]Re-running this example with $\mathbf{r} < 1$ nicely illustrates why we need Condition 3.

Chapter 7

Suppositional Consistency

Chapter 6 introduced Generalized Conditionalization (GC), CLF's updating constraint. We showed that (GC) allows CLF models to yield verdicts representing requirements of ideal rationality for stories involving memory loss or the threat thereof, stories that fall ourside the domain of a modeling framework based on the Bayesian's traditional Conditionalization updating rule. We then applied (GC) a couple of times to show that it yielded intuitive verdicts for a complex lottery story involving memory loss and a generalized version of van Fraassen's Reflection principle.

Still, it would be nice to have an intuitive understanding of the general requirement on doxastic evolutions represented by (GC). In this chapter I will argue that the basic idea behind (GC) is an idea that has lurked behind conditionalizing updating rules all along: what I call "suppositional consistency." Roughly speaking, suppositional consistency requires an agent to assign the same degree of belief to a claim whenever she considers it relative to the same conditions.

Explaining (GC) then becomes a project of explaining why suppositional consistency is a requirement of ideal rationality. It turns out there are two somewhat different answers one can give here, depending on whether one thinks ideal rationality specifies a unique required degree of belief for an agent relative to any body of evidence she might possess. If one believes this Credal Uniqueness thesis, there is good reason to think that evidential requirements have a structure mandating suppositional consistency. Whether they made this point explicit or not, a variety of historical figures who endorsed Credal Uniqueness offered substantive views that required suppositional consistency. A suppositional consistency requirement based on Credal Uniqueness also yields interpersonal requirements on degrees of belief that many authors have found plausible.

If on the other hand one denies Credal Uniqueness and admits that in some cases an agent's body of evidence might not mandate a particular degree of belief, it becomes unclear why the degrees of belief the agent actually assigns are required to line up over time at all. In Section 7.3 I will suggest that such diachronic requirements might come from doxastic commitments that take effect when an agent assigns a rationally permissible degree of belief. I will then argue that if there are such diachronic doxastic commitments, (GC) captures what they require of an agent's doxastic evolution. Still, there may be special circumstances in which diachronic commitments cease to apply; I will take up objections to (GC) in Section 7.4.

The upshot will be that whether a (GC)-based modeling framework applies to a particular story depends on a number of substantive theses in epistemology about the strength of ideal rationality's requirements, the existence of doxastic commitments, etc. Instead of arguing for a particular stance on each of these theses, this chapter tries to explain how (GC) would apply given each of the possible positions. Section 7.5 summarizes the discussion by describing how adopting (GC) affects CLF's domain of applicability.

7.1 The Basic Idea

Consider the following story:

> *Chocolate:* You and I are going to play a game. I will flip a fair coin; if it comes up heads, I will decide whether to give you a piece of chocolate. If the coin comes up tails, no chocolate for you. Before I flip the coin, you consider all your certainties about people in general, chocolate, and me in particular; of special relevance is your certainty that I'm someone who doesn't care very much for chocolate. As a result, you assign a 0.4 degree of belief that you will wind up with some chocolate. I then flip the coin and it comes up heads. Once you've seen this outcome, what does ideal rationality require of your degree of belief that you'll be receiving some chocolate?

A model of this story, model C, is described in Table 7.1. The first extrasystematic constraint applies the Principal Principle to this story, using your certainty at t_1 that the coin flip is fair. The last row in the table of extrasystematic constraints reflects the structure of the game; you are certain throughout that if the coin flip comes up tails there will be no chocolate for you.

Table 7.1: Model C

Story: Chocolate

T: Represents these times:

- t_1 After the game is explained to you but before I flip the coin.
- t_2 After the coin comes up heads but before I decide whether to give you some chocolate.

L: Built on these atomic sentences, representing these claims:

- h The coin comes up heads.
- c You receive some chocolate.

Extrasystematic constraints:

	P_1	P_2
h	$1/2$	1
c	0.4	< 1
$\sim h \supset \sim c$	1	1

Applying CLF to model C, (GC) yields the verdict

$$P_1(c \mid h) = P_2(c \mid \top) = P_2(c) \tag{7.1}$$

which we can then use to derive

$$P_1(c) < P_2(c) \tag{7.2}$$

Here model C is getting things right: Ideal rationality requires your degree of belief that you'll be receiving chocolate to increase when you learn the coin has come up heads. This will be a fairly direct consequence of any updating rule that directs you to conditionalize in response to a pure learning event. But can we say more about what precisely goes wrong when you violate such rules, or when you violate (GC) in general?

Since (GC) represents a requirement on conditional degrees of belief, it will help to review the positive account I gave in Section 5.3.2 of what goes on when an agent assigns a conditional degree of belief. I suggested that we can think of degree of belief assignment as a two-step process: first, the agent entertains some situation; second, she evaluates the claim in question in light of conditions in that situation. The situation is defined by what I will now call a **suppositional set**, the union of the claims of which the agent is certain and the claims she conditionally supposes in assigning her degree of belief. (For unconditional degree of belief assignments, the latter is empty and the suppositional set is just the agent's certainty set.) In Section 5.3.2 we discussed your degree of belief that you will go skiing next weekend conditional on the supposition that it will snow this week. In that example your suppositional set consists of your certainties about your skiing proclivities plus the additional supposition that it will snow this week; you then assign a confidence to the claim that you will go skiing relative to the conditions defined by that set.

We can think of this process another way, using slightly more philosophical machinery. Before I ask you to make any supposition you entertain a particular set of doxastic possibilities, each of which we can think of as a completely specified possible world. Your unconditional degree of belief that you will go skiing next weekend is an evaluation of that claim relative to this entire set of doxastically possible worlds. When I ask you to suppose that it will snow this week, you focus your attention on just those doxastic worlds in which it snows, and make an evaluation relative to that proper subset. But every claim of which you were initially certain still holds true in each of those worlds, since certainties hold across your entire doxastic possibility space.[1] Thus the set of worlds relative to which you now evaluate your conditional degree of belief is defined by the claims I ask you to suppose and the claims you already held certain. That is, it is defined by your suppositional set.[2]

Back now to Chocolate. You may or may not assign a degree of belief at t_1 to the claim that you will receive some chocolate conditional on the claim that the coin comes up heads, but for simplicity's sake let's assume that you do. When you make that assignment, you evaluate your prospects for chocolate in the situation in which the coin comes up heads—that is, you judge my inclination to decide in your favor should I be forced to decide. In doing so, you picture a situation that combines your t_1 certainties about how the game is played with a further supposition (whose truth is unknown to you at t_1) about the outcome of the flip.

Between t_1 and t_2 you see the outcome of the flip and become certain that you are living in the situation you imagined earlier. Your unconditional degree of belief that you will receive chocolate is now just your evaluation of your prospects for chocolate in the situation in which the coin comes up heads. If you assign this unconditional degree of belief a different value than you assigned conditionally at t_1, your evaluations of my inclination are *diachronically* inconsistent—you assign a different value to the same claim relative to the same set of conditions at two different times. To avoid this inconsistency, your degrees of belief must line up as prescribed in Equation (7.1).

The basic idea here is that an agent violates the requirements of ideal rationality if at two points she assigns the *same* claim relative to the *same* situation (whose conditions are defined by her suppositional set) *different* degrees of belief. Expressed as a formal constraint, this becomes

> **Suppositional Consistency:** For any sentence $x \in L$, times $t_j, t_k \in T$, and consistent set $S \subset L$ such that $C_j \subseteq S$ and $C_k \subseteq S$,
> $$P_j(x \,|\, \langle S - C_j \rangle) = P_k(x \,|\, \langle S - C_k \rangle)$$

Here S represents the suppositional set, which we require to be consistent so that the conditional credences receive defined values under the Ratio Formula. If the agent supposes $S - C_j$ at t_j and $S - C_k$ at t_k, the union of

[1](Easwaran 2008) proposes a probabilistic formalism for subjective degrees of belief that includes an underlying set Ω representing the agent's doxastic possibilities at a given time.

[2]For more on conditional degrees of belief—along with some helpful diagrams—see (Edgington 1996).

the claims she takes for certain and the claims she conditionally supposes will come to the same thing (that is, S) at each time. Suppositional consistency requires her to assign the same conditional degree of belief to an arbitrary claim (represented by x) relative to these identical suppositional sets.[3]

Theorem D.2 shows that given the synchronic framework, suppositional consistency is equivalent to (GC). In other words, adopting (GC) as CLF's updating rule is tantamount to adopting a requirement of suppositional consistency.[4] Taking suppositional consistency as the basic idea behind (GC), we can ask two questions: (1) Why should we believe that ideal rationality ever requires suppositional consistency, and (2) in what kinds of stories does it do so? The answers to these questions depend on the answers to broader questions about the strength of ideal rationality's requirements.

7.2 The Synchronic Solution

7.2.1 Credal Uniqueness and Conditional Structure

An agent is suppositionally consistent if she always assigns the same degree of belief to a claim relative to the same suppositional set. This will clearly be a requirement if, relative to any consistent suppositional set, there is some specified degree of belief that ideal rationality requires any agent at any time to assign a given claim. If there is a specified degree of belief ideal rationality requires for a claim relative to a set, the agent's assigning different degrees of belief to that claim relative to that set at different times means she's missed the specified value at least once, thereby violating the requirements of ideal rationality.

For example, consider the Chocolate story. At t_1 you have various certainties (about people, about chocolate, about me, etc.), and we can think of your certainty set at t_1 as representing your total evidence at that time. In epistemology in general it's probably not a good idea to equate an agent's

[3]Thinking in terms of possible worlds, suppositional consistency requires the agent to always assign the same claim the same degree of belief relative to the same set of doxastic possibilities. So, for example, when an agent undergoes a pure information-loss event, her set of doxastically possible worlds will expand, but according to suppositional consistency her doxastic evolution will violate the requirements of ideal rationality unless she maintains the same degree of belief relative to the set of doxastic possibilities she entertained before the event. Before the event this assignment was an unconditional degree of belief; after the information loss it is a degree of belief conditional on the conditions that pick the old set of doxastic possibilities out from her new expanded set.

[4]My understanding of suppositional consistency was greatly aided by (Levi 1980, Chapter 4) and a conversation with Michael Caie.

evidence with her certainties—a claim can be part of my evidence even if I'm slightly uncertain whether it's true. But in constructing stories for CLF analysis we set aside this uncertainty, treating claims in evidence as givens by making agents certain of them at various times. As we discussed in Section 3.1, this allows us to focus on how an agent's evidence influences her non-extreme degrees of belief. So for our purposes we will take certainty sets to represent an agent's total evidence at particular times.[5]

To this point we've been assuming that your t_1 degrees of belief in Chocolate—and in particular your $P_1(c) = 0.4$ assignment—do not already violate the requirements of ideal rationality. Now clearly there are some possible $P_1(c)$ assignments that would be inconsistent with your evidence; for example, assignments greater than or equal to $1/2$.[6] But perhaps given your evidence at t_1 there is exactly one rationally permissible $P_1(c)$ assignment, one precise degree of belief such that if you assign any other $P_1(c)$ value your doxastic evolution violates the requirements of ideal rationality. This would be consistent with a thesis White calls "Uniqueness":

> **Uniqueness:** Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition. (White 2005, p. 445)[7]

Uniqueness holds that an agent's total evidence mandates a particular doxastic attitude towards any proposition, but it does not say what kind of doxastic attitude is mandated. We can strengthen Uniqueness to a position I'll call "Credal Uniqueness," which holds that given any (consistent) total evidence set and any particular claim there is a precise *degree of belief* that ideal rationality requires an agent to assign to that claim when her total evidence is that set.[8] If Credal Uniqueness is true, 0.4 might be the wrong value for $P_1(c)$ given your total evidence at t_1; you might have already violated the requirements of ideal rationality when you assign that degree of belief.[9]

---

[5]If you're worried that claims in evidence must meet various conditions that certainties often don't, there's no reason we can't confine our attention to stories in which an agent's certainties are always *true*, or even *known*. (Compare (Williamson 2000, Ch. 9).)

[6]Following the standard interpretation's Certainty Conditions, I am assuming you assign $P_1(h \supset c) < 1$. Along with the first extrasystematic constraint on model C, this entails that your t_1 degree of belief in chocolate is required to be strictly less than $1/2$. (A similar argument requires it to be strictly greater than 0.)

[7]White, in turn, attributes the thesis to (Feldman 2007). I've quoted White's formulation of the thesis here.

[8]The term "Credal Uniqueness" is adapted from (Levi 1980, Section 4.2).

[9]Notice that the proposed requirement on your $P_1(c)$ value is a requirement *given your*

In Section 7.3 we will step back and discuss whether Credal Uniqueness is a reasonable position; for the rest of this section, I want to discuss what a Credal Uniqueness position might imply about the diachronic constraints on an agent's degrees of belief. Certainly a Credal Uniqueness view will require an agent to assign the same degree of belief to the same claim at any two times at which her total evidence sets are identical. But that's not enough to yield substantive required relations between the agent's degrees of belief across times when her evidence varies. To get that kind of structure, a Credal Uniqueness view needs to say something about the relations between the degrees of belief required on different total evidence sets.

Philosophers who hold Credal Uniqueness often do say something about such relations; they offer a substantive theory describing the degrees of belief required by ideal rationality in particular situations. Such a theory can be captured by what I will call a ***C*-function**: a function describing the degrees of belief ideal rationality requires an agent to assign to various claims given various bodies of evidence. In particular, if we are working with a time set T and a modeling language L, $C(h, E)$ will assign a real number to ordered pairs of $h \in L$ and $E \subseteq L$. $C(h, E)$ is the degree of belief the agent is required to assign the claim represented by h when her total evidence is represented by E. We can implement a C-function in a CLF model by going through each $t_i \in T$ and placing extrasystematic constraints on the model setting various $P_i(x)$ equal to the value of $C(x, C_i)$.[10]

In many C-functions a special role will be played by the values of $C(\cdot, \mathsf{T})$ (where T is a tautology in L). We can think of these as the degrees of belief it would be rational for an agent to assign who had no empirical evidence. This is a very old idea, running from the discussion of "*a priori* probablities" in (Keynes 1929) through the discussions of "initial probabilities" in (Carnap 1950) and (Jeffreys 1973) and up to the notion of a "reasonable initial credence function" in (Lewis 1980).[11] Many Credal Uniqueness adherents have held that the rationally required degrees of belief for an actual agent

total evidence, represented by the claims you take for certain at t_1. Thus the suggested violation of the requirements of ideal rationality still results from an *internal* inconsistency among doxastic attitudes within your doxastic state.

[10]On many substantive Credal Uniqueness theories, the degree of belief assigned to a particular claim relative to a particular total evidence set depends on the modeling language being used. That's why the implementation process starts by choosing a modeling language and then invokes a C-function. For more on the language dependence of such credence-assignment theories, see (Halpern and Koller 2004).

[11](Carnap 1955) describes the initial probability of a hypothesis as "its probability before any factual knowledge concerning the individuals is available." For references to more recent discussions of such "hypothetical priors," see (Meacham 2008, n. 7).

can be found by starting with the degrees of belief required of a hypothetical agent with *no* evidence and conditionalizing on the total evidence possessed by the actual agent. More specifically, Credal Uniqueness adherents have described C-functions such that for any T and L, the function displays

> **Conditional Structure:** Given a model with time set T and modeling language L, a function $C(h, E)$ defined over ordered pairs of sentences $h \in L$ and sets of sentences $E \subseteq L$ has conditional structure if it meets the following requirements:
>
> 1. $C(h, \mathsf{T})$ assigns a real number to each $h \in L$.
> 2. $C(\cdot, \mathsf{T})$ is a probability function (it satisfies Non-Negativity, Normality, and Finite Additivity).
> 3. For all $t_i \in T$, $C(\langle C_i \rangle, \mathsf{T}) > 0$.
> 4. For all $h \in L$ and $E \subseteq L$, if $C(\langle E \rangle, \mathsf{T}) > 0$ then $C(h, E) = C(h \,\&\, \langle E \rangle, \mathsf{T})/C(\langle E \rangle, \mathsf{T})$.

The key condition in the definition of Conditional Structure is the last: It expresses an arbitrary $C(h, E)$ value as a ratio of two "initial probabilities" and allows one to obtain $C(h, E)$ values by conditionalizing $C(\cdot, \mathsf{T})$ on one's evidence.[12]

Theorem D.5 shows that if you start with a model in the synchronic framework (that is, a model under the standard interpretation implementing only the synchronic systematic constraints of CLF) and then implement a C-function with conditional structure, the result will satisfy every requirement of (GC). Put another way, a conditional-structure C-function builds in a requirement of suppositional consistency.

This allows us to connect our discussion of suppositional consistency with historically important Credal Uniqueness positions. The *locus classicus* of Credal Uniqueness in the literature is (Carnap 1950). There Carnap describes how to construct a probability function $\mathfrak{m}^*(h)$ that assigns a positive value to every logically consistent proposition h. He then uses $\mathfrak{m}^*$ to construct a two-place function $\mathfrak{c}^*(h, E)$ which he offers as an explication of

[12]Conditional Structure puts a set of relations in place between $C(h, E)$ values for various Es an agent might actually have. It is possible to maintain these relations between required degrees of belief relative to realistic bodies of evidence without insisting that required degrees of belief relative to *no* evidence make sense. In fact, it is mathematically possible to institute such $C(h, E)$ relations for realistic Es without assigning any values to $C(\cdot, \mathsf{T})$. But historically most philosophers who have defended substantive $C(h, E)$ functions have employed an initial probability construction, so I have adopted that approach here.

"probability$_1$", his notion of evidential support.[13] Since $\mathfrak{m}^*$ is a probability function defined over all propositions, it satisfies conditions 1 and 2 above. Since every certainty set is consistent and $\mathfrak{m}^*$ assigns positive values to consistent propositions, $\mathfrak{m}^*$ satisfies condition 3 as well. Carnap constructs $\mathfrak{c}^*$ from $\mathfrak{m}^*$ exactly as described in condition 4 (with $\mathfrak{m}^*$ as $C(\cdot, \mathsf{T})$ and $\mathfrak{c}^*$ as $C(h, e)$). So $\mathfrak{c}^*$ has conditional structure, with $\mathfrak{m}^*$ playing the role of the initial probability function $C(\cdot, \mathsf{T})$.

Carnap holds that an agent whose total evidence is E should assign degree of belief $\mathfrak{c}^*(h, E)$ to the proposition h. Since $\mathfrak{c}^*$ has conditional structure, (GC) becomes a theorem of Carnap's system of inductive logic.[14] On Carnap's position, suppositional consistency is a universal requirement of ideal rationality.

Various authors have adopted Credal Uniqueness positions similar to this Carnapian one.[15] Jeffreys (1973, Chapter II) introduces a function $P(q \,|\, p)$ representing the reasonable degree of belief in p on data q, then provides axioms for this function that give it conditional structure. Similarly, Maher (2004) introduces a conditional-structure function $p(H \,|\, E)$ for the degree of belief in H that is justified by evidence E. Williamson (2000) begins with a function $P(p)$ (playing the role of $C(p, \mathsf{T})$), then carries out the construction described in condition 4 to build a function $P_\alpha(p)$ (playing the role of $C(h, \alpha)$) with conditional structure. For Williamson, $P_\alpha(p)$ represents the "evidential probability" of p on the evidence α. Its conditional structure allows Williamson to demonstrate that evidential probability updates by conditionalization "when evidence is cumulative" (2000, p. 220). In other words, Williamson starts with a C-function with conditional structure and derives Limited Conditionalization (LC)—an updating rule we saw in Chapter 6 is entailed by (GC).

Another popular current Credal Uniqueness view calculates its C-function by maximizing the entropy of an agent's distribution subject to the constraints provided by her evidence at a given time. (Seidenfeld 1986) explains this approach and shows that as long as the constraints an agent's degrees of belief are subject to at different times are consistent with each other, the

[13]Carnap represents an agent's total evidence with a "long sentence," (1950, p. 20) so his $\mathfrak{c}^*$ function actually takes a sentence as its second argument. In this discussion I've changed the second argument to a set to make connections between Carnap's approach and ours more clear; one could get all the results I describe here from Carnap's actual theory by using $\langle E \rangle$ as the evidential report sentence.

[14]Levi notes this fact (in his own, somewhat different terminology) at (Levi 1980, p. 86).

[15]The position should really be called "early Carnapian," since Carnap had abandoned it by (Carnap 1952). I'll continue to use the "Carnapian" label for simplicity's sake.

C-function provided by maximum entropy will have conditional structure.[16]

But *why* do these Carnapian views adopt a C-function with conditional structure—*why* do these authors set up C-functions that make suppositional consistency a requirement of ideal rationality? An agent's total evidence describes her current situation to the extent she understands it. According to Credal Uniqueness, whatever the agent takes her situation to be there is a unique degree of belief she is required to assign to any claim. But presumably it is the *content* of the claims in the agent's total evidence set—the *conditions* defining the situation for the agent—that determine that unique degree of belief. If at another time the agent *imagines* a situation meeting those conditions (without taking them all of them to be actual), the requirements relative to those conditions should still be the same. How a situation is arranged determines how confident of a claim the agent should be, whether the situation is taken as real or just supposed. So $C(h, E)$ describes not only the required degree of belief in h for an agent whose total evidence is E, but also the required degree of belief in h when the agent's suppositional set is E. Since there is a required degree of belief in any claim relative to any (consistent) suppositional set, if an agent assigns degrees of belief to the same claim relative to the same suppositional set at two times, those degrees of belief are required to be the same.

Having understood *why* Carnapian positions underwrite suppositional consistency (and therefore (GC)), it is important to see precisely *how* they wind up doing so. On Carnap's view, for example, an agent's t_j and t_k degrees of belief don't wind up conforming to (GC) because her t_j assignments rationally constrain her assignments at t_k. Instead, the agent's t_j and t_k degrees of belief are each synchronically constrained by $\mathfrak{c}^*$ in light of her total evidence at those respective times. The conditional structure of the $\mathfrak{c}^*$ function gives rise to a particular mathematical relationship between $P_j(x \mid \langle C_k - C_j \rangle)$ and $P_k(x \mid \langle C_j - C_k \rangle)$. If that relationship does not hold between the agent's t_j and t_k degrees of belief, the agent's assignments at at least one of those times differ from what $\mathfrak{c}^*$ requires. In a sense, the Carnapian does not understand (GC) as a *diachronic* constraint at all, nor suppositional consistency as a truly diachronic requirement. According to the Carnapian, when an agent's assignments fail to fit the pattern described by (GC) her doxastic evolution violates the requirements of ideal rationality because her degrees of belief at some time are out of step with the evidence

[16]As Formal (GC) makes clear (Section 6.1.3), (GC) does not yield verdicts when an agent's certainty sets at two times are inconsistent with each other. So the condition on Seidenfeld's result should not keep the maximum entropy enthusiast from endorsing (GC).

embodied by her certainties *at that same time*. Notice that this kind of argument can be made even if one doesn't know precisely what values the C-function assigns to various (h, E) pairs. As long as one knows that there is a C-function and that it has conditional structure, one can be guaranteed that an agent whose doxastic evolution violates (GC) is violating a synchronic requirement of ideal rationality somewhere.

We can think of Credal Uniqueness positions as affirming the existence of a *global* C-function; a C-function that assigns a required degree of belief to any claim relative to any (consistent) evidential set, and so will specify all the credence values in a model for any T and L you choose. A great number of philosophers will deny the existence of a global C-function, because they reject Credal Uniqueness as a general thesis. Yet while they deny that every consistent total evidence set mandates a unique degree of belief for every claim, these philosophers will nevertheless grant that some evidence sets do mandate degrees of belief for some claims.[17] So they will grant that some stories are covered by a *local* C-function, a function that requires unique degrees of belief for the particular claims and particular bodies of evidence that happen to be of interest in that story. And relative to the T and L used in modeling the story, that local C-function may have conditional structure.

For example, someone who denies Credal Uniqueness in general may still endorse the Principal Principle, which requires agents with particular bodies of evidence to set their credences in line with objective chances. In a story in which *all* the degrees of belief of interest are governed by the Principal Principle (for example, The Die from Chapter 3), the relevant chance function will generate a local C-function. And while I won't defend this claim here,[18] a local C-function generated entirely by objective chance values will have conditional structure. Thus Principal Principle adherents who reject Credal Uniqueness will nevertheless accept that (GC) indicates genuine requirements of ideal rationality in such "chance-governed" stories.

Some philosophers reject Credal Uniqueness because it is an evidentialist thesis: it maintains that rationally required doxastic attitudes supervene on the agent's evidential state.[19] Opponents of evidentialism may hold that the degree of belief required of an agent with a particular total evidence set depends on various other contextual matters; Levi, for example, lists "[the agent's] goals and values, the problems he is investigating, the way

[17]It's untenable to hold that no evidence set mandates a degree of belief for *any* claim—after all, a consistent evidence set mandates extreme degrees of belief for the claims it deductively entails or refutes.

[18]The details needed for a proof can be found in (Lewis 1980).

[19]For evidentialism, see (Conee and Feldman 2004).

he has succeeded in identifying potential solutions, and other circumstantial factors." (1980, p. 92) Still, these philosophers may grant that in stories throughout which the non-evidential factors remain constant, a local C-function with conditional structure is available, making (GC) applicable once more.

The point here is not to give an exhaustive catalog of all the available non-Carnapian positions. The point is that while the Carnapian position takes the requirements of ideal rationality to be particularly strong—strong enough to dictate a unique required degree of belief for every agent in every evidential situation—positions that take the requirements of ideal rationality to be weaker may nevertheless be able to make a synchronic argument for suppositional consistency in a wide variety of stories. Even if we do not know what precise degrees of belief ideal rationality requires an agent to assign in a particular story, we may be able to convince ourselves that a local C-function with conditional structure is available. In that case (GC) will yield verdicts indicating genuine requirements of ideal rationality and the story will fall within CLF's domain of applicability. On the other hand, if we *are* supporters of Credal Uniqueness, both the historical record of Credal Uniqueness positions and the general Uniqueness view of evidential relations strongly suggest that the universal C-function has conditional structure. This makes (GC) a theorem, and places a wide variety of stories within CLF's domain of applicability.

7.2.2 Conditional Structure and Interpersonal Relations

Suppose we are analyzing a story for which the requirements of ideal rationality provide a C-function with conditional structure. The C-function may be local and apply only to the bodies of evidence encountered in this story, or (if we believe in Credal Uniqueness) it may be the universal C-function governing every evidential state conceivable. Either way, the agent's degrees of belief in such a story will be required to line up in the manner described by (GC). But as we have seen, this requirement derives from the underlying mathematical structure of the C-function and the synchronic constraints that function puts on the relations between the agent's degrees of belief at a given time. The fact that the diachronic degrees of belief being related belong to the same agent is incidental.

Thus in stories governed by a conditional-structure C-function, (GC) can represent required relations among degrees of belief belonging to different *agents* just as well as relations among degrees of belief belonging to the same agent at different times. In Section 4.3.3 we mentioned an interpretation of

CLF in which the members of the time set represent not times but agent-time pairs: t_1 might represent agent A at noon, t_2 might represent agent A at 1pm, and t_3 might represent agent B at 1pm. Verdicts of a CLF model relating $P_1(x)$ and $P_2(x)$ would indicate required relations between agent A's degrees of belief in the same claim at different times, while verdicts relating $P_2(x)$ and $P_3(x)$ would indicate required relations between the degrees of belief assigned to the same claim by agents A and B at the same time.

Keeping the rest of the standard interpretation roughly intact, (GC) would represent a requirement of ideal rationality for stories in which all the agents' degrees of belief were governed by a C-function with conditional structure. Our talk of Jill's and Ken's degrees of belief—which was introduced as a metaphor in Section 6.1.3 to help explain (GC)—now becomes literal: conditional on the claims Ken is certain of but she isn't, Jill is required to assign the same degrees of belief as Ken does conditional on all the claims Jill is certain of but he isn't. Using Formal (GC) instead of the "official" version, we can put this another way: Suppose C_j represents Jill's certainty set and C_k represents Ken's. If $C_j \cup C_k$ is consistent, then relative to that suppositional set Jill and Ken must assign the same degrees of belief.

In Section 6.2.2 we saw that under particular conditions, (GC) and the standard interpretation yield a Generalized Reflection principle governing an agent's degrees of belief conditional on suppositions about her degrees of belief at other times. Under the interpersonal interpretation of CLF, we can derive a principle governing an agent's degrees of belief conditional on suppositions about *other* agents' degrees of belief. I won't work through the details, but here's the result: Suppose we let P_a represent agent A's degrees of belief and P_b represent agent B's. Under particular conditions (similar to the conditions enumerated in Section 6.2.2), ideal rationality requires

$$P_a(x \,|\, [P_b(x \,|\, \langle C_a - C_b \rangle) = \mathbf{r}]) = \mathbf{r} \tag{7.3}$$

Put into words, when agent A supposes that agent B assigns degree of belief $\mathbf{r}$ to claim x conditional on all the certainties A possesses but B doesn't, ideal rationality requires A to assign $\mathbf{r}$ to x as well. Elga (2007) recommends this principle for dealing with a "guru"—someone whose judgment you trust and who possesses information you don't, but who may also fail to possess some information you do. (For example, a weather forecaster who knows much more than you about general weather patterns but doesn't know that it's raining outside your window right now.) And just as Generalized Reflection entails Reflection in situations in which the agent's certainty set strictly increases, the interpersonal Generalized Reflection principle implies what

Elga calls an "expert" principle[20] for relating to agents who possess all of your information and more:

$$P_a(x \,|\, [P_b(x) = \mathbf{r}]) = \mathbf{r} \tag{7.4}$$

In this case, conditional on the supposition that the expert (agent B) assigns unconditional degree of belief $\mathbf{r}$ to x, ideal rationality requires agent A to assign $\mathbf{r}$ to x as well.[21]

7.3 Doxastic Commitments

There is a sense in which much of our defense of a (GC)-based modeling framework to this point has involved cheating. The applications of (GC) we described in Chapter 6 involved stories (The Die, Shangri La, The Lottery) in which the requirements on the agent's degrees of belief at any time could (arguably) be derived synchronically from her evidence at that time and the Principal Principle. Thus in these stories a conditional-structure C-function—grounded in an objective chance function—was available to govern the agent's degrees of belief. Given the results of Section 7.2.1, it is no surprise that these applications supported the (GC)-based framework.[22]

If we believe in Credal Uniqueness, this is no problem. On this view, there is a universal C-function describing the degree of belief in any claim required for an agent with any body of evidence. As we argued earlier, this

[20] Elga takes the "expert" terminology from (Gaifman 1988).

[21] Early in his (1991), Christensen writes that the Reflection principle requires us to "regard our own future selves quite differently (epistemologically speaking) from the way we view other people." (p. 232) Later (p. 245), he presents the expert principle we've just described under the name "Solidarity" and criticizes it as "wacky" when applied in general. But we are defending the principle only in contexts in which the agent is certain that the expert is at least as knowledgeable as she and satisfies the requirements of ideal rationality, and in which the story in question is covered by a C-function with conditional structure. In such contexts, there is no important epistemological difference between the doxastic attitudes of the expert and the (ideally rational) doxastic attitudes of our future selves. Far from being wacky, interpersonal Reflection-style principles seem quite reasonable in such cases.

[22] I am hardly the only one guilty of this cheat. Authors often curry favor for a proposed conditionalization updating rule by applying it to stories in which all of the agent's degrees of belief are dictated by some norm like the Principal Principle that provides a conditional-structure C-function. The updating rule does a beautiful job of yielding verdicts that match our intuitions, but that's because all of the structure needed to generate those verdicts arises synchronically. If an updating rule works only in such situations, there is a sense in which it is only a delightful mathematical pattern someone has noticed and not an independent diachronic *constraint* in its own right.

C-function will have conditional structure and will mandate degrees of belief over time that satisfy (GC).

But suppose we deny Credal Uniqueness as a general thesis. This seems a fairly reasonable position; after all, in many cases we think two agents with the same evidence can adopt different doxastic attitudes towards a claim without either one's violating the requirements of ideal rationality. And it can seem like a strech to think that one's evidence, no matter how vague or sketchy, always determines a unique degree of belief in any claim (down to an arbitrary number of decimal places!) required by ideal rationality. We may admit that some stories are covered by C-functions, and we may even think that whenever a C-function is available it has conditional structure. But the pressing question is whether CLF yields correct verdicts for stories for which no C-function is available. Why should an agent's evolving degrees of belief over time be required to conform to (GC) in cases in which particular degrees of belief aren't required of her to begin with?

For example, let's return to the Chocolate story. As we already noted, ideal rationality requires your degree of belief in the prospect of chocolate before the coin is flipped to be between 0 and 1/2. Since you are certain at t_1 that I don't care very much for chocolate, perhaps your degree of belief should be higher than 1/4. But does that limited information about my chocolate proclivities really dictate a precise numerical degree of belief that ideal rationality requires you to assign to the prospect of chocolate? And if there is no specific $P_1(c)$ value required, why should $P_2(c)$ be required to relate to $P_1(c)$ in the manner dictated by (GC)? Put another way, if there is no particular precise judgment that ideal rationality requires you to make at t_1 about my inclination to decide in your favor should the coin come up heads, why should it require you to make the same judgment at t_2 as you made at t_1?

7.3.1 Denying Credal Uniqueness

Let's suppose that in the Chocolate story your evidence and the requirements of ideal rationality do not mandate a precise t_1 degree of belief that you will receive some chocolate. In other words, let's assume that Credal Uniqueness is false and Chocolate provides a counter-example.

Some epistemologists hold that an agent should not adopt a doxastic attitude towards a claim unless that attitude is required by her evidence.[23] Coupled with the plausible thesis that a body of evidence can require at most

[23] Compare principle (b) at (Levi 1980, p. 89) and White's discussion at (White 2005, p. 447).

one doxastic attitude towards a particular claim, this position runs into a problem if it admits that there are stories (like Chocolate) for which Credal Uniqueness fails. That problem can be solved, however, by proposing that in Chocolate you should adopt a t_1 doxastic attitude towards the prospect of chocolate that is best represented by a credence *range* $(\mathbf{a}, \mathbf{b})$. (Such "ranged attitudes" were introduced in Section 4.3.2.) We might suggest that while not every evidence set dictates a precise degree of belief required of the agent for every claim, every evidence set does dictate a particular ranged attitude for each claim.[24] Notice that this view is a version of White's Uniqueness thesis—that thesis never required the "unique rational doxastic attitude" dictated by one's total evidence to be a precise degree of belief. The view also makes a doxastic attitude required by her evidence available to an agent in every evidential situation while denying Credal Uniqueness for stories like Chocolate in which the thesis seems implausible.

CLF has been designed to allow one to work with credence ranges if one so desires. As we noted in Section 4.3.2, there is an alternative interpretation of CLF that allows its models to yield verdicts indicating requirements of ideal rationality on doxastic evolutions involving ranged attitudes. On that interpretation (GC) represents a requirement of ideal rationality on how an agent's *ranged* attitudes evolve over time. While I won't work through the details here, (GC) still turns out to be equivalent to a notion of suppositional consistency, where that notion involves an agent's assigning the same *ranged attitude* to the same claim relative to the same suppositional set at any two times. Moreover, we can define "ranged" C-functions that take an ordered sentence-set pair (h, E) to a range $(\mathbf{a}, \mathbf{b})$, and develop a notion of conditional structure on which ranged C-functions with conditional structure give rise to requirements of suppositional consistency. So stories covered by ranged C-functions with conditional structure will still have (GC) as a diachronic constraint.

Yet if the motivation for Ranged Attitude Uniqueness (as we might call it) is supposed to be that ranged attitudes provide plausible uniquely-required doxastic attitudes in every story while precise degrees of belief do not, I don't think it succeeds much better than Credal Uniqueness. Why should we think that one's evidence, no matter how vague or sketchy, always determines a unique required ranged attitude (with bounds to an arbitrary number of decimal places!) in any claim? One tempting answer is that at a given time an agent's evidence will rule *out* particular degrees of belief in a claim, and ideal rationality requires the agent to adopt the ranged attitude

[24]Christensen (2004, p. 149), for example, considers this suggestion.

represented by all the values not ruled out. For example, in Chocolate the combination of the Principal Principle with CLF's synchronic constraints and Certainty Conditions rule out $P_1(c)$ values greater than or equal to $1/2$ and equal to 0 (for an explanation see note 6 above). Perhaps ideal rationality requires you at t_1 to adopt the ranged attitude towards your prospects of chocolate represented by $(0, 1/2)$. But this range does not take into account your evidence about my aversion to chocolate. To dramatize the point, suppose there is an initial time t_0 in the story when the game has been explained to you but you are not yet aware that I don't care for chocolate. By the reasoning above you should assign the $(0, 1/2)$ ranged attitude to your prospects of chocolate at t_0 as well. Yet surely ideal rationality requires some change in your attitude towards the claim that you will be receiving chocolate when you learn that I—the person who has a good chance of making a decision about whether or not to give it to you—don't care much for chocolate.[25]

My point here is just that Ranged Attitude Uniqueness does not seem much more plausible than Credal Uniqueness. Many philosophers favor ranged attitude representations of doxastic attitudes for very sensible reasons, for example a desire to respect the Keynesian distinction between uncertainty and risk,[26] or a sense that numerical ranges provide a more faithful representation of agents' actual psychology in some situations than single-value degrees of belief.[27] But these motivations need not lead to the dubious claim that precise required numerical ranged attitudes are always available while precise required numerical degrees of belief are not.[28]

[25]A Ranged Attitude Uniqueness defender might respond that your attitude towards the chocolate claim should remain identical (that is, at $(0, 1/2)$) from t_0 and t_1, it's just that your attitudes towards *other* claims (such as "he doesn't like chocolate" or "I have a high chance of getting some chocolate") should change in light of your new evidence. Yet it seems awfully odd to deny that this evidence should have any *direct* influence on your attitude towards the claim that you'll be receiving chocolate. Moreover, that position threatens to eviscerate the Bayesian account of relevance, according to which one claim is relevant to another if learning the former rationally alters an agent's attitude towards the latter. Surely "He doesn't care for chocolate" is *relevant* to "I will receive some chocolate."

By the way, I think this argument is equally effective against positions according to which you should adopt a special doxastic attitude at t_1 towards your prospects for chocolate described as "withholding judgment," or adopt no doxastic attitude towards that claim at that time at all. Both these positions will have to say the same thing about your attitude towards that claim at t_0, and then deny that there should be a change in your attitude between the two times.

[26]As in (Weatherson 2002).

[27](Christensen 2004, p. 149, note 4) has a brief but interesting discussion of this point.

[28]See (Walley 1996, p. 10) on this point.

Given the apparent implausibility of both Credal Uniqueness and Ranged Attitude Uniqueness, we might disavow the Uniqueness thesis altogether. White describes views that deny Uniqueness as "epistemically permissive." An epistemically permissive view need not be anything-goes; for Chocolate an epistemically permissive view might concede that degrees of belief equal to or above 1/2 are forbidden, while permitting any degree of belief assignment that remains. Or a less permissive epistemically permissive view could require that in light of your evidence about my chocolate proclivities only degrees of belief above 1/4 are allowed. What makes a view epistemically permissive is that it allows situations in which there are multiple doxastic attitudes an agent could take towards a particular claim without violating the requirements of ideal rationality. Notice that this means denying the claim that a doxastic attitude is *permitted* by ideal rationality only if it is *required*.

7.3.2 Making a Commitment

For the rest of this chapter we will adopt an epistemically permissive view and focus our attention on cases in which ideal rationality permits different degree of belief assignments to a claim relative to the same body of evidence. (To streamline discussion we will also set aside consideration of ranged attitudes.) We will ask why there should be any diachronic constraints on an agent's doxastic evolution at all in such situations, and if there are why they should be of the kind represented by (GC).

Taking up the first question, we can think of an agent who assigns a degree of belief in an epistemically permissive situation as making a sort of judgment call in selecting among the degrees of belief rationally permitted to her. We may also think that such a judgment call comes with a set of doxastic commitments. For example, an agent who assigns a high degree of belief to a claim is committed to assigning a low degree of belief to its negation.[29] Finally, we may think that an agent who assigns a degree of belief that is permitted but not required has not only synchronic but also diachronic commitments: she is committed to responding to possible future pieces of evidence by altering her degrees of belief in particular ways. If this is the case, then (GC) can be read as an attempt to represent the *content* of an agent's diachronic doxastic commitments. If we have set up CLF, its intended interpretation, and its systematic constraints correctly, a CLF

[29] Or, more precisely, to not assigning anything other than a low degree of belief to its negation (compare our discussion of certainties and commitments in Section 4.2.2).

model will indicate a violation of the requirements of ideal rationality when an agent fails to honor her diachronic doxastic commitments.[30]

Suppose, for example, that the Chocolate story involves an epistemically permissive situation, and that your $P_1(c) = 0.4$ assignment is permitted but not required. Taking this assignment into account, model C yields verdicts that $P_1(c \mid h) = 0.8$ and therefore $P_2(c) = 0.8$. In assigning $P_1(c) = 0.4$, you are committed to assigning no degree of belief other than 0.8 to the prospect of chocolate should you learn that the coin came up heads; if you assign an unconditional t_2 degree of belief to chocolate other than 0.8, you renege on a doxastic commitment and violate the requirements of ideal rationality.

If this doxastic commitments view is right, the requirements of ideal rationality represented in CLF enforce doxastic commitments in epistemically permissive situations, even though the agent was not required by ideal rationality to undertake those particular doxastic commitments to begin with. (That is, the agent could have assigned a different initial degree of belief and thereby undertaken a different doxastic commitment without violating the requirements of ideal rationality.) Yet one might wonder how an initial degree of belief assignment can place requirements on an agent's later assignments if that initial assignment was not required by ideal rationality to begin with. If it would have been perfectly permissible for the agent to make a different assignment at an earlier time, why isn't it permissible for the agent not to honor that assignment at a later time?

People who press this objection on me often go on to suggest that our feeling that something has gone wrong if you assign, say $P_1(c) = 0.4$ but $P_2(c) = 0.3$ is a hangover from the general feeling that there must be specific correct values for $P_1(c)$ and $P_2(c)$, and that the doxastic evolution under consideration gets at least one of these values wrong. (In other words, the objector thinks that diachronic commitment defenders are being driven by an implicit conviction that suppositional consistency can be given sychronic

[30]It's tempting to say that whenever an agent makes a judgment call she *puts* various diachronic commitments in place. But I will try to avoid talking that way, because I don't want to take a stand on whether a judgment call puts diachronic commitments in place when it violates the requirements of ideal rationality. (For instance, if you assigned $P_1(c) = 0.75$ in the Chocolate story.) The precise thing to say is that according to CLF's diachronic constraints an agent who doesn't line up her later degrees of belief with her earlier assignments in the manner represented by (GC) has a doxastic evolution that violates the requirements of ideal rationality. (Notice that this will be trivially true in the case in which the earlier assignment stands in violation on its own.) We may think of this in terms of diachronic commitments put in place by the earlier assignment, or we may think of it in terms of a general, standing commitment to diachronic consistency. (Thanks to Nico Silins for discussion on this point.)

grounds.) Certainly an agent who assigns a degree of belief in an epistemically permissive situation often thinks it's the only rationally permissible assignment, and so will see an error if other people (or even her own self at other times) assign degrees of belief that fail to match. But if you truly accept (the objector argues) that a $P_1(c) = 0.4$ assignment is rationally permissible in light of your evidence at t_1, and that a $P_2(c) = 0.3$ is rationally permissible in light of your evidence at t_2, then there is no reason a doxastic evolution combining those two assignments shouldn't be rationally permissible as well. Ideal rationality should not require an agent to honor a degree of belief assignment that was rationally arbitrary to begin with.[31]

It's odd that no one ever presses this objection in the synchronic case. If your evidence at t_1 permits an assignment of $P_1(c) = 0.4$ and permits an assignment of $P_1(\sim c) = 0.7$, why aren't you permitted to assign both these doxastic states at once? Why should ideal rationality require your $P_1(\sim c)$ assignment to honor a $P_1(c)$ assignment that is rationally arbitrary to begin with? To the extent that epistemologists adopt epistemically permissive positions, they do not see this as a reason to disavow synchronic requirements such as those represented in the probability axioms. This suggests to me that the argument form "Such-and-such doxastic attitude is not rationally required, so it cannot place such-and-such rational constraints on such-and-such other doxastic attitudes" is invalid.

But even having shed this argument *against* the existence of diachronic doxastic commitments, it would be nice to be able to say something in their favor. Here I think the best strategy is to strip the case down to its bare essentials. The following story presents the most straightforward diachronic commitments case I can think of:

> *Baseball:* The A's are playing the Giants tonight, and Ray and Ken are discussing who will win the game. They agree that it's a tough matchup to call: the Giants have better pitching, but the A's have a more potent offense; the A's have won most of the matchups in the past, but the A's are weaker this year than usual. All in all, a rational person could go either way. Nevertheless Ken asks Ray what he thinks, and Ray says "I'm not certain either way, but I'm leaning towards the A's."
>
> Five minutes later, Bill comes in and asks Ray who he thinks will win tonight's game. Ray says, "I'm not certain either way, but I'm leaning towards the Giants."

[31] White, for example, presses this sort of objection at (White 2005, pp. 454-5).

This series of responses strikes us as puzzling and in need of explanation. Following the grand Davidsonian tradition, we might try to explain how these responses could have made sense from Ray's point of view.[32] We might suggest that Ray gained some new relevant information between answering the two questions; perhaps he glanced through his A's media guide and saw a crucial statistic he wasn't aware of before. Perhaps Ray remembered a relevant fact about the matchup that he hadn't thought about for a long time and wasn't taking into account when he provided his initial answer. Perhaps Ray's responses don't really reveal his beliefs, and there's some pragmatic reason why he would give a different response to Bill than to Ken. Perhaps the very sight of Bill puts Ray in a contrarian mood and influences his judgment.

Presumably there are questions we could ask Ray to test these various explanations. Suppose we ask him those questions and it turns out that none of them is an accurate description of his experience between giving his two answers. We concoct some more explanations and ask him about those, but they aren't correct either. In the end, Ray admits to us that he just believed one thing at one time and another thing at another. I think that in this case we would conclude that Ray's series of beliefs was irrational. Even in an acknowledged epistemically permissive situation, there is a rational failing in a diachronically inconsistent doxastic evolution.

We might draw an analogy here to an issue in the theory of practical reason. Consider an example from (Bratman 1987, p. 23). Bratman is driving from Stanford to San Francisco and has equal reason to take Highway 280 or Highway 101—ideal rationality will not be violated if he takes either route. He forms an (admittedly arbitrary) intention to take Highway 101, and to do so he must turn right at Page Mill Road. But suppose that when he reaches Page Mill Road he does not turn right and takes 280 instead. Has he violated any requirements of ideal rationality? (After all, the 101 route was never required by ideal rationality to begin with!)

This example nicely focuses our attention on what rational *consistency* requires of an agent. In the example there are no external factors (so to speak) that make 101 better for Bratman to take when he reaches the crucial turnoff. But since he formed an earlier intention to take 101, it would be *inconsistent* of him to drive off in the other direction. Similarly, at t_2 in Chocolate your evidence doesn't favor one $P_2(c)$ assignment over another. But given the judgment you made at t_1 about my inclination to give you

[32]Davidson provides a nice summary of his approach in the opening pages of (Davidson 1982).

chocolate should the coin come up heads, it seems *inconsistent* for you to assign a $P_2(c)$ value other than 0.8.

7.3.3 The Structure of Doxastic Commitments

I certainly have not presented a knock-down case in favor of diachronic doxastic commitments; it is still very open to an epistemic permissivist to deny that there are such things. But for the time being let's put aside such denials. Let's also put aside anti-evidentialist positions on which pragmatic factors that change over time can influence the requirements on an agent's degrees of belief—or at least let's confine our attention to stories in which those factors are held constant. When we consider permissive situations in which diachronic doxastic commitments are in force and only epistemic factors exert an influence, I think there's a good case to be made that (GC) (in partnership with the other elements of CLF) accurately captures the structure of those diachronic commitments.

We can think of diachronic consistency as a kind of constancy. The relevant constancy is not constancy of one's unconditional degrees of belief in claims. It is our plight as doxastic agents to receive evidence in a gradual trickle rather than all at once upon birth, and this inflow of evidence should change our confidence in claims over time. But we can maintain constancy in how we evaluate claims relative to *situations* of the sort described in our account of conditional degrees of belief. As time goes on we learn that some of these situations are non-actual, but we can maintain any degrees of belief in claims we assigned relative to situations that remain live possibilities. This is the requirement of suppositional consistency, which we have seen is equivalent to the requirement represented in (GC).

Let's start with some simple examples. If diachronic consistency requires an agent at a later time to respect assignments she made at earlier times, surely she is required to assign the same degree of belief to a claim at two times at which her evidence is the same. Similarly, she is required to assign the same degree of belief to one claim conditional on another claim at any two times at which her evidence is the same. The only thing suppositional consistency adds is a requirement of constancy between two points at which the union of what the agent is supposing and what she takes for certain is the same, but the dividing line between suppositions and certainties falls in a different place. For example, suppositional consistency requires an agent whose certainty set is $\{c, d\}$ to assign the same value to $P(a \mid b)$ as she assigned to $P(a \mid b \,\&\, c)$ when her certainty set was just $\{d\}$.

Notice that in a synchronic setting this dividing line makes no difference to an agent's conditional degrees of belief. Consider all the credences at a particular time t_i conditional on $\langle S \rangle$ (for some set $S \subseteq L$). We could record those in a function $P_{i/S}(x)$ defined over $x \in L$, where $P_{i/S}(x) = P_i(x \,|\, \langle S \rangle)$. Assuming $P_i(\langle S \rangle) > 0$, our synchronic systematic constraints guarantee that $P_{i/S}(\cdot)$ is itself a probability function (see Theorem A.13). And that function makes no distinction between the members of S and the members of C_i. Members of each receive a $P_{i/S}$ value of 1, and they influence the other $P_{i/S}(x)$ values in exactly the same way.

In other words, when an agent assigns a conditional degree of belief at a given time it makes no difference which claims fall on which side of the boundary between conditional and supposed. This suggests that if between two times an agent keeps her suppositional set constant but simply moves the dividing line between certain and supposed, consistency requires her to keep her degrees of belief the same.

Further support for this conclusion is provided by CLF's successes with epistemically permissive stories. We have already seen that a CLF model yields verdicts representing requirements of ideal rationality for the Chocolate story. But we can also create epistemically permissive versions of The Die, Shangri La, and The Lottery. For example, suppose that in The Lottery (Section 6.2.1) Dave has no idea whether the contestants' names will be drawn from the hat using a fair process. Conspiracy theories have been swirling around him, but he decides to discount them all and assign an initial 1/10 degree of belief to each contestant's prospects. If Dave operates under a set of diachronic doxastic commitments, I think we will agree that ideal rationality requires his later degrees of belief to develop exactly as the CLF model we described for that story indicated. Although The Lottery as we first described it was covered by a C-function, I don't think our judgments about the diachronic requirements of ideal rationality in that story ultimately depend on the agent's initial degree of belief assignments' being traceable to objective chances. The same point could be made about The Die and Shangri La.

Once we grant that ideal rationality requires an agent to honor a set of doxastic commitments over time, there is good reason to believe that suppositional consistency characterizes those commitments. This conclusion is supported by our understanding of the nature of conditional degrees of belief, by comparison to synchronic constraints on conditional degrees of belief, and by the correct verdicts we have seen CLF yield for a number of stories over the course of this chapter and the last. Even Levi, who thinks that an agent's doxastic commitments can be voided by shifting non-evidential contextual factors, agrees that when they hold those commitments have the mathematical structure represented in (GC). Because Levi is working within a ranged attitudes framework his presentation is rather different from ours, but once we shift to a ranged attitudes interpretation of CLF, Levi's "confirmational tenacity" (1980, Chapter 4) becomes equivalent to our requirement of suppositional consistency.

7.4 Objections to (GC)

I now want to respond to four objections to (GC). The first two strike me as objections that could be raised by someone who believes in diachronic doxastic commitments but doubts (GC)'s account of them; the last two attack doxastic commitments as we're understanding them here. (Keep in mind that we're still working under the assumption that Credal Uniqueness is false.)

7.4.1 Unique Updates

Critics of Bayesianism often complain that Conditionalization-style updating rules require an agent to assign an initial degree of belief to every claim she entertains conditional on every piece of evidence she might receive, so that she is fully prepared to update by conditionalizing no matter what evidence comes her way. This is sometimes put by saying that a Bayesian model can apply only to an agent who starts off (from birth?) with a full "prior" distribution over all the claims she will ever entertain.[33] Yet this does not follow from (GC) on CLF's standard interpretation. Under the standard interpretation's Evaluative Rule, (GC) requires an agent's final degrees of belief to match up with initial conditional degrees of belief in particular ways *if she assigned those initial conditional degrees of belief.* As we saw in Section 4.2.2, a verdict will not indicate a violation of the requirements of ideal rationality if the agent fails at the revelant times to assign degrees of belief to claims represented in the verdict.

So, for instance, if (GC) yields a verdict of the form $P_j(x \,|\, \langle C_k - C_j \rangle) = P_k(x \,|\, \langle C_j - C_k \rangle)$ for a model, this will indicate that *if* the agent assigns

[33]cf. (Harman 1986, Chapter 2): "One can use conditionalization to get a new probability for P only if one has already assigned a prior probability not only to E but to $P \,\&\, E$. If one is to be prepared for various possible conditionalizations, then for every proposition P one wants to update, one must already have assigned probabilities to various conjunctions of P together with one or more of the possible evidence propositions and/or their denials. Unhappily, this leads to a combinatorial explosion.... Doing extensive updating by conditionalization... would be too complicated in practice." Compare also Earman's discussion of his (LO2) at (Earman 1992, p. 122ff.).

a particular t_j conditional degree of belief and a particular t_k conditional degree of belief, and *if* these are not equal, *then* her doxastic evolution violates the requirements of ideal rationality. An agent needs to assign a wide range of conditional degrees of belief at an initial time only if she wants to make sure her final degrees of belief will be highly constrained by what she believed earlier. But surely it does not violate the requirements of ideal rationality for an agent to take up at a later time claims she did not think through earlier, or combinations of claims and evidence she did not consider at an earlier time.[34] Since the agent did not pursue these lines earlier, when she finally takes them up her degrees of belief may be fairly unconstrained diachronically. But that simply means that upon taking up new claims for consideration or considering old claims in the light of not-fully-thought-through evidential combinations, the agent must now make up her mind on matters she never got around to making her mind up about before.[35]

Yet we may have a concern in the other direction. One of the attractions of conditionalization updating rules has always been the promise that they would uniquely dictate how an agent should respond to receiving new information. If an agent updates her degrees of belief on a particular occasion by conditionalizing, then given a full initial degree of belief distribution and the set of certainties upon which she conditionalizes, we can derive all of her final unconditional degrees of belief. Viewed from the agent's point of view, conditionalization can provide a *recipe* for changing one's degrees of belief in a rational fashion in response to new information.

As we've just seen, an agent rarely has a full initial degree of belief distribution across all the evidence sets she might encounter, so the vision of an updating rule that would provide a unique rational response to any possible evidential situation was always a bit of a pipe dream. Nevertheless, it may be seen as a disadvantage of (GC) that it will not always yield a unique required final credence distribution even given an initial distribution over an entire modeling language. For pure learning events, (GC) will take a complete initial distribution and yield a complete final distribution. But when a different kind of doxastic event occurs between two times, (GC) will

[34] These were the sorts of examples we examined in Section 5.2.

[35] Of course, even if the agent doesn't assign the particular conditional t_j degree of belief represented in a (GC) verdict, its value may be rationally constrained by her other t_j assignments and this may in turn constrain the conditional P_k value. For example, in Chocolate your unconditional assignments to $P_1(c)$, $P_1(h)$, and the like will constrain $P_1(c \,|\, h)$ and therefore $P_2(c)$, even if you don't actually assign a t_1 degree of belief to your prospects of chocolate conditional on the coin's coming up heads.

yield only constraints on final conditional credences, not precise values for all final unconditional credences (unless some credences from before the two times in question provide further constraints on those, as in the transition from t_1 to t_2 in Shangri La).[36]

Of course, if the story in question is covered by a C-function, that function will create enough synchronic requirements to pick up the slack and require unique final unconditional degrees of belief for every claim represented in the modeling language. But suppose we are working with an epistemically permissive situation. Is there always a unique degree of belief assignment that ideal rationality requires after, say, a memory-loss event?

For some time now philosophers and artificial intelligence researchers have studied various "logics of belief revision," including most prominently the Alchourrón-Gärdenfors-Makinson (AGM) model. (Alchourrón, Gärdenfors, and Makinson 1985) These models describe how an agent's entire belief set should adjust when the agent changes her beliefs about a particular claim. In particular, the AGM model yields a unique belief set an agent should adopt after a "contraction" in which the agent ceases to believe a particular claim.[37]

AGM was designed to model an agent's full beliefs, not necessarily her certainties, and the sorts of belief contraction envisioned result from defeaters like discovering that a belief was based on unreliable evidence.[38] Still, we can imagine a belief revision model that describes how an agent's certainty set should change when she loses information due to memory loss. We can further imagine a model (perhaps an AGM-Bayesian hybrid?) that

---

[36]A number of people have suggested something like this objection to me; I am grateful to Greg Restall and Aidan Lyon for putting it in particularly clear and forceful ways.

[37]Those familiar with the AGM literature may notice similarities between (GC) and AGM's "recovery" postulate—especially in the way (GC) requires an agent to return to her original doxastic state if she loses a certainty and then regains it—and may be concerned that CLF will fall prey to well-known counterexamples to recovery. (I am grateful to Graham Priest for suggesting this point and to Gordian Haas for further discussion.) In examples in which an agent loses a certainty then regains it, all the changes to her certainty set (in particular making her final certainty set identical to her initial certainty set) are managed by stipulations in the story and so by machinery outside CLF. CLF begins its modeling once the agent's evolving certainty set has been fully described, stepping in to represent patterns in the agent's non-extreme degrees of belief that follow from her changing certainties. The standard counterexamples to AGM would therefore operate at the extrasystematic level on which certainty sets are determined, and so constitute objections to the stories themselves rather than to the analysis of these stories by CLF. I do not know if parallel examples can be constructed that generate problems for CLF itself; despite a bit of trying I have not managed to concoct any. (For a nice summary of AGM and some counterexamples to recovery, see (Hansson 2006).)

[38]We will discuss the effects of such discoveries on an agent's certainties in Chapter 12.

tells us not only how the agent's certainties should change when she forgets, but also what her new degrees of belief should be in the claims she has forgotten.

Yet I doubt such a model is possible. The fundamental problem is that the stories such a model would aim to analyze are underspecified. If, for example, we have a pure information-loss story that describes the agent's initial degree of belief distribution and then tells us which certainties are lost, there is not enough information in the story to dictate a unique required final degree of belief distribution. Since no unique final distribution is required, no model can tell us what the required final distribution should be.

In Chapter 6 we suggested that from CLF's formal point of view, memory loss is like learning backwards in time. Suppose I describe an agent's current full degree of belief distribution over a set of claims, tell you which of those claims she has learned since some particular earlier time, then ask you for her full degree of belief distribution at that earlier time.[39] You can't produce that earlier distribution, but the problem isn't you: the problem is that the story I've told doesn't provide enough information to fully specify the agent's initial unconditional degrees of belief. To provide that information, the story would have to tell you not only which claims the agent learned since the earlier time, but what her degrees of belief in those claims were before she learned them.

The same goes for memory-loss events. We are imagining that an agent is certain of a particular claim at a given time, then loses that certainty due to memory-loss at a later time. We are further imagining that this is an epistemically permissive situation, so that no requirements of ideal rationality based strictly on her evidence at the later time mandate a precise required degree of belief in the claim in question. Finally, we are imagining that the agent does not retain any doxastic commitments from an even earlier time (as in the relation between t_0 and t_2 in Shangri La) that would force her to a particular required degree of belief. (If there were such lingering commitments, (GC) would pick them up.) What the agent has to do, then, is make an entirely fresh judgment call at the later time as to what her unconditional degree of belief in the claim will be. The outcome of this judgment call is not mandated by the requirements of ideal rationality, so there is no way that a modeling framework designed to represent the requirements of ideal rationality can predict what it will be. Once we are told what judgment is made, we can represent it in our model and test it for consistency with the

[39] Assume that the set of claims I describe her current distribution over does not include second-order claims describing her earlier assignments.

rational constraints that do apply to the case. But within the space that is epistemically permitted by those constraints, there is no way the framework can specify a unique outcome.

Bayesians have gotten used to modeling stories involving only pure learning events. Their experience with this special case has led them to expect that given an initial distribution and a description of the changing certainty sets, an updating rule will yield a full, unique required final distribution. But when we widen the domain of applicability of our modeling framework to include pure information-loss and mixed doxastic events, we no longer obtain such unique solutions. My suggestion is that the fault lies not in our framework, but in the information provided. The story as described underdetermines the agent's required final distribution, so it's no surprise that CLF cannot tell us what that final distribution should be.[40]

7.4.2 Reevaluating Evidence

One way to describe the requirement of suppositional consistency is as a requirement that one keep fixed one's evaluations of various claims in light of various bundles of evidence, actual and potential. But shouldn't new evidence be able to affect how one evaluates claims in the light of evidence? To take an example, suppose I assign conditional degrees of belief to some claims with respect to various suppositional sets that include evidence from testimony. Then a friend of mine comes along and asserts a claim to which I had previously assigned a very low unconditional degree of belief. After hearing this bit of testimony, I presumably should adjust my degree of belief in that claim. But shouldn't I also downgrade my estimate of the reliability of testimony, even if only slightly? And doesn't this suggest that I should now evaluate claims (including the claim asserted) relative to bundles of evidence including testimony differently than I did before?[41]

There is a sense in which gaining new evidence changes how we view other evidence. For example, at the beginning of a hand of seven-card stud I might not view evidence that my last card is going to be the Jack of Hearts

[40]An analogy: Imagine a physics student who has only ever modeled elastic collisions (using something like our framework NMF). He decides one day to model inelastic collisions, and complains that there must be something wrong with the laws conserving kinetic energy and momentum because they cannot take information about masses and velocities before the collision and generate unique velocities for after. The student's limited experience has led him to expect more from the laws than they can be asked to give; if the scenario he is modeling does not specify the energy lost to heat in the collision, it simply does not determine unique velocities for after the collision.

[41]I am grateful to Stanley Chen for conversations on this point.

as evidence that I'm going to win the hand. But once my first three cards turn out to be the other three jacks, any evidence that my last card is going to be the Jack of Hearts becomes incredibly good evidence that I'm going to win the hand. In this case my views of what counts as evidence for what have rationally changed as I gained more evidence about the contents of my hand.

Our unconditional degrees of belief change as evidence comes in (or is forgotten), and so do our conditional degrees of belief. This means that when we consider one piece of evidence in isolation, we may over time change our views on whether that piece of evidence should boost or reduce our confidence in a particular claim. But that's very different from changing our views on how the claim should be evaluated in light of a *total evidence set*, the kind of thing that defines an imagined situation. Before I drew any cards, I could have imagined the situation in which I had the three other Jacks and evaluated how important the Jack of Hearts would seem to me then; this evaluation shouldn't change when I actually have the three Jacks in front of me. So the cards case is no counterexample to suppositional consistency. (It better not be, since it is covered by a very simple conditional-structure C-function generated by the objective chances of drawing various cards from the deck!)

Now let's return to the testimony example. Before my friend makes his assertion, I may assign a conditional degree of belief to the asserted claim conditional on the supposition that he asserts it. I picture a situation in which my friend makes that assertion, and evaluate the claim in light of conditions in that situation. One of the conditions in that situation is that my friend has made a seemingly outrageous assertion, and if I do a thorough, thoughtful job of evaluating the claim in light of those conditions part of that evaluation will be an estimation of their implications for the reliability of testimony. That is, whatever adjustments I make to my appraisal of testimony when I actually hear my friend's assertion, I should make the same adjustments when assigning a degree of belief conditional on the supposition that my friend will make that assertion. All the adjustments a piece of evidence requires in my lowest-level degrees of belief, my degrees of belief about the relevance of evidence to those lowest-level degrees of belief, etc. can and should be made the same way when I conditionally suppose a piece of evidence as when I acquire that evidence for real.

Of course, we have all had the experience of pondering a possible eventuality in advance, thinking we understood its implications, then when that eventuality actually comes into being reconsidering what we thought before. But in those cases I think we will admit that there is a flaw in our doxastic evolution: either our earlier doxastic state was flawed because we weren't thinking the eventuality through all the way,[42] didn't assign proper weight to what it would actually be like, etc.; or our later doxastic state is flawed because we are letting the experience overwhelm us, unduly influence our judgment, etc.; or both (as I suspect often occurs). If at the earlier time we were truly assigning a degree of belief conditional on *all* the information we gain between the two times, and if we are assigning our degrees of belief now in a level-headed, doxastically responsible fashion, we should evaluate the same claim the same way on both occasions.

7.4.3 Changing Your Mind

Go back to Ray in the epistemically permissive Baseball story from Section 7.3.2. When Ray first backs the A's and then backs the Giants, we might ask him incredulously, "Didn't you just say the opposite five minutes ago?" What if Ray responds that until you brought it up he'd forgotten that he did so? Or what if Ray says that in the intervening five minutes he simply changed his mind about who's going to win the game?

Let's start with mind-changing. We are imagining a situation that is truly epistemically permissive, in the sense that two agents with the same evidence may make different judgment calls in assigning degrees of belief to claims without either one's violating the requirements of ideal rationality. An agent changes her mind in one of these situations if she makes one judgment call at an earlier time, then reconsiders and makes a different one later on. We do this kind of thing all the time, and it doesn't seem irrational. For example (pursuing our practical reason analogy from earlier), Bratman might between the moment he decides to take Highway 101 and the moment he makes the turnoff change his mind and decide to take Highway 280 instead. Yet the suppositional consistency requirement we have been describing seems to forbid changing one's mind: if an agent assigns one degree of belief to a claim relative to a particular suppositional set at one time and a different one later, her doxastic evolution is suppositionally inconsistent even if she has "changed her mind" in the interim.[43]

It isn't quite right to say that suppositional consistency (or a (GC)-based modeling framework) *forbids* an agent from changing her mind. As

[42]Some failures to think things through may be failures of logical omniscience—failures to see that one claim entails another, say—but they need not all be. We may simply have missed something that we would have recognized as an evidentially relevant connection had we seen it, even though the relation involved is not a strictly deductive one.

[43]I am grateful to Matthew Parrott for pressing this point.

we said in Chapter 4, the standard interpretation's Evaluative Rule provides an *evaluative* standard, not a set of prescriptions. We have never said that an agent *should* always adopt that doxastic state that would prevent her doxastic evolution from violating any requirements of ideal rationality. Nevertheless, it is true that by CLF's lights the doxastic evolution of an agent who changes her mind in the manner described above violates the requirements of ideal rationality. Notice that this is an evaluation of the doxastic evolution, not of the agent who changes her mind.

While we may therefore go easy on mind-chang*ers*, CLF takes a critical stance towards mind-chang*ing*. I will assume that the mind-changing under discussion is something one *does*, and does consciously. ("You know, I was thinking a bit more about what my broker said and changed my mind about the riskiness of that investment. . . .") If in the Baseball example Ray just had one doxastic attitude at one moment and then had a conflicting one a bit later without having devoted any more thought to the matter, we wouldn't describe him as having changed his mind—we'd just say he'd been irrational.

It might be suggested that changing one's mind must violate the requirements of ideal rationality, because it is doxastically irresponsible—it is a change in attitude not driven by a change in one's evidence. But the stories we are considering are ones that we have already conceded to be epistemically permissive, meaning that we have already permitted agents in these stories to assign degrees of belief that are not wholly driven by evidential considerations. Why apply stricter standards to an agent when she changes her mind than we do when she makes it up in the first place?

Mind-changing is a strange and slightly mysterious epistemic action, and I think the question of its rationality will ultimately be decided by a much deeper theory of the nature of doxastic commitments than what I am attempting here. (Perhaps that theory will involve the constitution of epistemic agenthood and the continuity of one's identity as an agent over time. . . .) But just as a start, notice that what we have been calling a "judgment call" need not be made on a whim—just because it is evidentially arbitrary does not mean it is *entirely* arbitrary, or even entirely *epistemically* arbitrary. Different agents may have different epistemic values that drive their degree of belief assignments in permissive situations. For instance, one scientist may consistently prefer the simplest of the available hypotheses consistent with her evidence, while another may choose the hypothesis with the most predictive power. Even if neither preference is forced by ideal rationality (or by the scientists' evidence), it may be that for a scientist committed enough to one view of the scientific enterprise it would be irrational to form the opinions suggested by another. For the prediction-preferer to change her mind away from an attitude driven by that preference would be inconsistent with her deep epistemic outlook and with a number of assignments made on other occasions.

On the other hand, some judgment calls—like Ray's about the baseball game—may just be spur-of-the-moment reactions not rooted in any epistemic values that could coordinate opinions over time. Perhaps when an agent makes one of these judgment calls, it is rationally permissible for her to remake them as many times as she sees fit. It may be that for mind changes shallow enough not to involve seriously changing one's *mind*, (GC) holds agents to a more restrictive standard than ideal rationality does.

7.4.4 Forgetting an Earlier Assignment

These considerations will also help with the case in which Ray forgets his earlier degree of belief assignment. Our attitude in this chapter and the preceding one has been that an agent does not violate the requirements of ideal rationality simply by forgetting a certainty. In the Path by the Sea Shangri La case, you are certain of the claim that you travel the Path by the Sea while you are traveling that path, then come to be less-than-certain of that claim once you reach Shangri La as a result of the guardians' memory tampering. I have argued that this is all consistent with the diachronic requirements of ideal rationality. But your earlier certainty about which path you travel was itself a doxastic attitude, and presumably it came with diachronic commitments of its own. The reason you aren't required to honor those commitments once you reach Shangri La (that is, the reason you aren't required to go on being certain you traveled the Path by the Sea) is that you have not only lost the certainty, but also no longer remember which certainty it was that you had. The principle seems to be that if you no longer remember the doxastic attitude that anchors a doxastic commitment, ideal rationality no longer requires you to honor that commitment either.[44]

So is Ray's picking the Giants after the A's permissible if in between he forgot that he had earlier picked the A's? Here we have to be careful

[44]For someone who holds this principle it may be tempting to replace diachronic constraints (in permissive cases) with synchronic constraints describing what an agent should think given what she thinks she thought earlier. While I would endorse some such synchronic constraints (such as those that can be derived from Generalized Reflection), I have a hard time understanding how they can be motivated unless we have the diachronic constraints as well. Matching what you think were your earlier assignments makes sense to me only as an *attempt* to match actual earlier assignments that ideal rationality requires you to honor.

about what it is to forget an earlier doxastic attitude. Certainly we can have degrees of belief in higher-order claims about our lower-order attitudes. But our earlier judgment calls are not retained solely as explicit attitudes towards claims of the form "I used to assign such-and-such conditional degrees of belief." When an agent makes up her mind about a matter, her decision may leave doxastic traces that remain long after she loses track of what she thought about matter in particular. Or, as we discussed in the previous section, a judgment call may be driven by deeper epistemic values that linger after the original matter is completely forgotten.

As with mind-changing, this may be a case in which the requirements of ideal rationality on an agent who makes a judgment call depend on exactly what type of call is made and what underlies it. If a degree of belief assigned in an epistemically permissive situation is driven by a set of values that linger, the agent may be required to honor those values and make new assignments matching the old even if the explicit earlier assignments have been forgotten. For shallow, more whim-like judgment calls, forgetting the earlier call may entitle you to make a new one.

Yet we should be wary of letting agents off the hook for forgotten commitments, because it creates the possibility of a rational dilemma already nicely illustrated by Ray's case. Suppose we adopt the position that agents are required to be suppositionally consistent with all and only the earlier assignments that they remember. (Taking into account that remembering an early assignment need not involve an explicit attitude towards the claim that one made that assignment earlier on.) Once Ray is reminded that he picked the A's a few minutes before he picked the Giants, he is in an impossible situation: his earlier A's judgment commits him one way, his Giants judgment commits him the other, and there is now no way for him to honor all the doxastic commitments he remembers.[45] Moreover, this is a pretty realistic doxastic situation—we've all had the experience of forgetting a piece of information, being completely unable to retrieve it at one time, and then suddenly remembering it later!

One could offer various responses to this scenario—perhaps the requirements of ideal rationality can conflict, perhaps they can lead to no-win situations, etc. Rather than chase the dialectic further, I will simply propose the

---

[45]This is somewhat like a situation in which an agent confronts two experts who recommend different doxastic attitudes towards the same claim. Our discussion of expert and guru principles in Section 7.2.2 was in the context of stories covered by C-functions; in such a story if two experts disagree about the import of a body of total evidence then at least one of them is irrational. But in a permissive story conflicts can arise between perfectly rational experts with the same total evidence.

following constraint on CLF's domain of applicability: CLF verdicts derived from (GC) cannot be trusted to represent requirements of ideal rationality in epistemically permissive stories in which an agent at some point ceases to be bound by diachronic commitments as a result of forgetting non-extreme doxastic attitude assignments. This restriction takes care of all the troublesome forgetting cases I have been able to construct so far. I leave it to further epistemological investigation to determine the exact circumstances in which forgotten assignments void diachronic doxastic commitments.[46]

7.5 (GC) and CLF's Domain of Applicability

Chapter 6 began by pointing out a serious limitation on the Conditionalization-based framework's domain of applicability. A framework whose updating constraint is Conditionalization will indicate a violation of the requirements of ideal rationality whenever an agent goes from certainty to less-than-certainty in a claim. Yet in at least some forgetting stories (such as Shangri La) a loss of certainty does not violate the requirements of ideal rationality. So the Conditionalization-based framework gets those stories wrong.

Chapter 6 then presented a new updating rule—Generlized Conditionalization (GC)—and showed that a (GC)-based framework indicates correct requirements of ideal rationality in stories the Conditionalization-based framework got wrong. So moving from Conditionalization to (GC) expands our framework's domain of applicability. The question is, how far? What new categories of story does a (GC)-based framework get wrong?

I find it hard to pinpoint the exact boundaries of the (GC)-based framework's domain of applicability. This is not because I have difficulties understanding (GC); (GC) is a fairly simple formal constraint and given any story I can tell what a (GC)-based model will indicate is required of the agent. The trouble is that I don't adequately understand what ideal rationality requires in some of these stories, so I have a hard time telling if (GC) is getting it right.

My hope in this chapter has been to indicate the kinds of questions in epistemology on which the boundaries of CLF's domain of applicability depend. For instance, a great deal rides on Credal Uniqueness (or Ranged Attitude Uniqueness, if we are working with an interpretation of CLF that allows for ranged attitudes). If Credal Uniqueness is true—if there is a unique degree of belief required of every agent for every claim in every

[46]Compare the interesting practical question of whether an agent is wrong when she fails to honor a promise she's forgotten that she made.

situation—then I think we have argued convincingly that the C-function expressing ideal rationality's requirements will have conditional structure. In that case (GC) will never yield verdicts that do not indicate true requirements of ideal rationality; adopting (GC) as CLF's updating constraint will not have limited CLF's domain of applicability at all.

If Credal Uniqueness is false, there will still be some stories that are covered by a local C-function dictating unique degrees of belief for the agents at every time in that story. Section 7.2.1 gave us good reason to believe that those C-functions will have conditional structure, and so that those stories will fall within CLF's domain of applicability.

But if Credal Uniqueness is false, there will also be stories not entirely covered by C-functions. These are the "epistemically permissive" stories in which multiple degree of belief assignments to the same claim relative to the same evidence comply with the requirements of ideal rationality. If one denies that there are such things as diachronic doxastic commitments at all, these permissive stories will fall outside CLF's domain of applicability, because (GC) will indicate diachronic requirements where there are none.

Still, I think the mainstream position among permissivists is that there can be genuinely diachronic requirements of rational consistency even when the evidence underdetermines an agent's attitudes. I believe we have provided strong arguments that (GC) captures the content of those commitments. The exceptions to suppositional consistency I described near the end of this chapter—raw mind changes and forgetting earlier non-extreme assignments—seem to be cases in which the doxastic commitments themselves are voided, not cases in which (GC) misdescribes them. Nevertheless, if there are such exceptions stories involving them will be outside CLF's domain of applicability.

If we ever discover the truth about such issues as evidentialism, Credal Uniqueness, and the existence and nature of doxastic commitments, the typology of possible stories will simplify and the story about CLF's domain of applicability will become more clean. It's important to note, though, that on virtually all of the possible positions a (GC)-based framework fares better than one based on the traditional Conditionalization updating rule. The restrictions we've mentioned on CLF's domain of applicability apply to the Conditionalization-based framework as well, but that framework lacks CLF's ability to yield substantive, accurate verdicts for stories involving memory loss.

We began this chapter with an idea that is basic to the entire concept of conditionalizing updates: the idea of suppositional consistency. It seems to me that conditionalization rules gained their initial intuitive appeal from the idea that if an agent contemplates the same situation at two times, she ought to assign the same degree of belief to a given claim. Yet once we focus on this idea, it becomes apparent that suppositional consistency can provide us with rational constraints for pure information-loss events and mixed doxastic events just as well as it can provide us with constraints for pure learning events. Those constraints are neatly represented in CLF's Generalized Conditionalization rule.