

## Time-slice epistemology and action under indeterminacy

Sarah Moss

ssmoss@umich.edu

There is a question that unifies several recent debates in epistemology, namely whether there are any essentially diachronic norms of rationality, or whether all fundamental norms of rationality are temporally local. Let us say that fans of temporally local norms advocate *time-slice epistemology*, where at a first pass, we define this theory as the combination of two claims. The first claim: what is rationally permissible or obligatory for you at some time is entirely determined by what mental states you are in at that time. This supervenience claim governs facts about the rationality of your actions, as well as the rationality of your full beliefs and your degreed belief states. The second claim: the fundamental facts about rationality are exhausted by these temporally local facts. There may be some fact about whether you are a rational person, for instance. But this fact is a derivative fact, one that just depends on whether your actions and opinions at various times are rational for you at those times.

Suppose that perdurantism is correct, i.e. that objects have temporal parts located at different times, just as they have spatial parts located in different places. Then we can restate time-slice epistemology as a theory that concerns time slices, or instantaneous temporal parts of objects. The first claim: what is rationally permissible or obligatory for a time slice is entirely determined by the mental states of that time slice. The second claim: the fundamental facts about rationality are exhausted by facts about the rationality of time slices. There may also be derivative facts about whether temporally extended agents are rational. For instance, we could say that you are rational just in case you are composed only of rational time slices, or just in case

---

1. Thank you to audiences at the University of Michigan and the 2013 Columbia-NYU Graduate Conference in Philosophy for helpful discussion of this paper. I am also especially grateful to Tom Dougherty, Brian Hedden, Susanna Rinard, Scott Sturgeon, Brian Weatherson, Robbie Williams, and an anonymous referee for extensive comments on earlier drafts.

most of your time slices are rational. The point is just that rationality is not fundamentally predicated of you, but of your time slices. In a nutshell: time slices are the fundamental subjects of epistemic evaluation.<sup>2</sup>

Time-slice epistemology may initially seem to have plenty of counterintuitive consequences. Say you see your friend Alice eat four scoops of ice cream for lunch, and after lunch you form the belief that she has not eaten anything all day. This seems like a perfectly good example of an irrational belief. Say you yourself eat seven scoops of ice cream for lunch, even though you are going to regret your binge as soon as it is over. This seems like a perfectly good example of an irrational action. In both cases, it is tempting to say that you are irrational precisely because there is no connection between your past mental states and what you currently believe, or between your future mental states and what you are currently doing. Hence these cases may seem like counterexamples to time-slice epistemology.

But the cases are not counterexamples. The time-slice epistemologist agrees that you are irrational in these cases, just not that you are irrational in virtue of ignoring what you used to believe or what you will later desire. The idea behind the theory is that you are irrational in virtue of ignoring your *current* beliefs and desires. For instance, you currently remember seeing Alice eat ice cream, and that is why it is irrational for you to believe that she has not eaten anything. You currently care about whether you are happy later, which is why it is irrational for you to do something that you believe will make you unhappy later. In more generality, the idea behind time-slice epistemology is that the current normative import of your past and future mental states is entirely mediated by your current mental states.

This paper is programmatic in nature, with two very general goals. First, I want to tie together several epistemological theories by identifying them as theories that advance time-slice epistemology. This goal is addressed in the first section of the paper, where I define and motivate time-slice epistemology.

Second, I want to suggest that analogies with ethical claims can help us defend certain time-slice theories, namely time-slice theories of action under indeterminacy. In §2, I discuss several theories about how you should act when the outcome of your decision depends on some indeterminate claim. I start with Caprice, a theory of action under indeterminacy defended in WILLIAMS 2013. Caprice says how agents should act in isolated, one-off decision situations. Caprice follows from a more complete theory, Liberal, that also says how agents should act when they face multiple decision

---

2. The initial statement of time-slice epistemology should make it clear that the theory is not committed to perdurantism. In the context of this paper, claims about time slices are merely convenient shorthand for claims about agents that instantiate certain properties at times, such as being in particular mental states or being permitted to do particular actions.

situations over time. Liberal is a time-slice theory. In §3, I defend it against objections. In particular, I defend Liberal against general objections to time-slice theories by comparing it with compelling ethical claims.

Although Caprice and Liberal are compelling, they are not perfect. In §4, I raise some objections to these theories. In light of these objections, I develop alternative theories of action under indeterminacy in §5. Here again, I rely on useful analogies with ethical claims as I develop more robust principles in support of time-slice epistemology.

## 1 Defining and motivating time-slice epistemology

Time-slice epistemology replaces diachronic norms of rationality with synchronic norms, norms that say what is rational for individual time slices when they have particular mental states. The most obvious target of the time-slice epistemologist is the classic updating norm introduced by BAYES 1763 and widely adopted by Bayesians, namely the claim that agents should update their credences according to Conditionalization. Conditionalization demands that your later credence in a proposition should match your earlier conditional credence in that proposition, conditional on any information you have since learned. The norm is at odds with time-slice epistemology, as it says your current credences are rationally constrained by your past credences, on which your current mental states need not supervene.

As time-slice epistemologists, our case against Conditionalization begins with the observation that, as WILLIAMSON 2000 put it, “forgetting is not irrational; it is just unfortunate” (219). There may be meaningful epistemic norms that require your memory to be perfect. But these are not norms of rationality in the ordinary sense. If you forgot what you ate for dinner last night, we might criticize you—but not by saying, ‘how *irrational* of you!’ In the context of an argument about whether Alice is a rational person, it is not obviously relevant to mention that she is forgetful. Intuitively, rational requirements on evidence retention are more similar to requirements on evidence gathering. Being *negligent* about what you learn or remember may signal or constitute irrationality. This is especially the case for strategic negligence, e.g. if you selectively forget or fail to gather evidence that disconfirms your favorite theory. But just as you are not irrational merely for having imperfect powers of evidence gathering, you are not irrational for having imperfect powers of evidence retention.<sup>3</sup>

In more generality, the problem is that on our traditional understanding of Conditionalization, the norm requires that evidence is cumulative for rational agents,

---

<sup>3</sup> I am grateful to an anonymous referee for suggesting this argument. See TALBOTT 1991 for a classic discussion of Conditionalization and memory loss.

whereas intuitively, rationality does not impose such strict demands. In fact, sometimes it imposes contrary demands. In the Shangri La case in ARNTZENIUS 2003, an agent is rationally required to violate Conditionalization, rather than remaining certain of a proposition for which she lacks sufficient evidence.

There are two very different strategies for responding to these challenges for Conditionalization. The first strategy is to simply modify the diachronic norm with some restrictions. TITELBAUM 2013 suggests this approach when he says that “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” (124). In fact, some theorists claim that this restriction is already implicitly understood in traditional discussions of the updating norm. For instance, SCHERVISH et al. 2004 respond to the examples in ARNTZENIUS 2003 by complaining that certain “restrictions or limitations” are “already assumed as familiar” when Conditionalization is applied, and that these restrictions include the constraint that agents not lose information over time (316).

The second strategy for responding to challenges for Conditionalization is to trade in this diachronic norm for a synchronic norm that will yield its intuitive verdicts, but also yield the right verdicts about cases of memory loss. For instance, WILLIAMSON 2000 argues that “a theory of evidential probability can keep separate track of evidence and still preserve much of the Bayesian framework” (220). According to Williamson, your current credences are not constrained by your past credences, but by your current evidence. At any given time, your current credence in a proposition should match the prior conditional probability of that proposition, conditional on your current evidence. The prior probability distribution is a distinguished measure of “something like the intrinsic plausibility of hypotheses prior to investigation” (211), and your current evidence is just your current knowledge (185). Since knowledge is not necessarily cumulative for rational agents, this proposal answers challenges involving rational memory loss. Since your current mental states include your current knowledge, this proposal advances time-slice epistemology.

The same strategies can be used to respond to other challenges to Conditionalization. For instance, it is a familiar observation that without forgetting anything, rational agents can start out certain that some *de se* proposition is false, and then later have some credence in that same proposition. For example: you may rationally be sure that it is not yet after midnight, and then later have some credence that it is after midnight. This sort of rational credal change is incompatible with Conditionalization as it is traditionally stated. Again, there are two strategies for responding to the challenge. Some theorists simply restrict the diachronic norm. For instance, TITELBAUM

2013 endorses “Limited Conditionalization,” which “looks exactly like Conditionalization, except that it applies only when [an agent] retains all certainties at the later time that she had at the earlier time” (124). By contrast, time-slice epistemologists will again trade in Conditionalization for a synchronic norm, one that yields the right verdicts about cases of *de se* updating. For instance, the theory of updating defended in Moss 2012 constrains your current credences in *de se* propositions using only facts about your current mental states, namely your current memories and your current opinions about the passage of time. To sum up so far: in response to counterexamples to diachronic norms, we can simply restrict those norms so that they do not apply in cases where they would yield bad results. Or we can come up with alternative synchronic norms that yield intuitive verdicts in the challenging cases.

All else being equal, the second strategy wins. From the point of view of theory building, the repeated restriction of diachronic norms is unsatisfying. The simplest diachronic norm would say that your present opinions should match your past opinions. But this norm fails when you have more information than your past self, or less information than your past self, or different *de se* information than your past self. The problematic cases for diachronic norms are exactly those cases where your past opinions do not have their usual effects on your current mental states. Usually you remember and trust what your past self believed. But when this connection fails, diachronic norms yield counterintuitive consequences. Time-slice epistemology is a natural response to this pattern of observations. Instead of restricting diachronic norms to cases where your past credences have their usual effects on your current mental states, we should admit that your current mental states are what determine whether your current credences are rational.

The same goes for the normative import of your future opinions and desires. The simplest future-oriented diachronic norm for belief would say that your present opinions should match your future opinions. This norm is false when you are not sure what you will later believe, when you are wrong about what you will believe, when you think you might get misleading evidence, or when you think you might later be irrational. The simplest diachronic norm for desire would say that your present desires should match your future desires. This norm fails when you are not sure what you will later desire, when you are wrong about what you will desire, or when you think you might come to have some despicable desires, either as a result of some rational change like joining another political party or some irrational change like being influenced by drugs. Usually you anticipate and trust what your future self believes, and you want what your future self desires. But when these connections fail, it is no longer intuitive to suppose that your future states should constrain your current

beliefs and desires.

The foregoing discussion reveals two motivations for preferring time-slice theories over restricted diachronic norms. First, time-slice theories are generally stronger, yielding verdicts in cases where restricted norms are silent. Second, time-slice theories are generally simpler. They do not include ad hoc maneuvers around potential counterexamples. Instead they describe the indirect ways in which your past and future mental states do affect what you should currently believe and desire. To appreciate the contrast, consider analogous debates about interpersonal norms of rationality. The simplest interpersonal norm for belief would say that your opinions should match the opinions of your neighbor. This norm fails when you are not sure what your neighbor believes, when you are wrong about what they believe, when you think that your neighbor has misleading evidence, or when you think that your neighbor is irrational. In response, we could restrict interpersonal norms so that they do not apply in these cases. But it is far more natural to say that how your neighbor affects what you should believe depends on how much you trust that she has reasonable beliefs. The normative import of her beliefs is mediated by your opinion about her. The time-slice epistemologist makes just the same move in the intrapersonal case.

This reflection on interpersonal norms offers more than just a helpful analogy. In fact, it suggests a second definition of time-slice epistemology. Time-slice epistemologists say that what is currently rational for you does not fundamentally depend on your past and future attitudes. This insight can be developed in multiple ways. First, what is rational for you might depend on your current mental states. As we have seen so far, we can define time-slice epistemology as the claim that the fundamental norms of rationality are *synchronic*, grounding the rationality of your current states and actions in facts about your current mental states. Second, what is rational for you might depend on very general relations that hold between your current self and your past and future selves. In particular, these relations might be general enough that they hold not just between distinct temporal parts of a single agent, but between distinct agents. Hence we could alternatively define time-slice epistemology as the claim that the fundamental norms of rationality are *impersonal*, so that your rationality is grounded in facts about normative relations that hold between persons as well as between temporal parts of persons. The handle 'time-slice epistemology' is better suited for the first theory, but not entirely unfit for the second. Compare: the United States Supreme Court's ruling in *Citizens United v. Federal Election Commission* confers meaningful legal status on corporations by conceiving of them as subjects of legal norms traditionally reserved for persons. The second notion of time-slice epistemology confers meaningful epistemic status on temporal parts of persons by conceiving

of them as subjects of epistemic norms traditionally reserved for persons. By featuring in fundamental epistemic norms, time slices play a more significant role in our theorizing about rational belief and action.

Like our first notion, our second notion of time-slice epistemology has been defended in recent literature. For instance, CHRISTENSEN 1991 observes that distinct agents are not rationally required to have beliefs that cohere with each other simply because they are guaranteed to lose money otherwise, and he concludes that the same goes for distinct time slices of individuals: “the guaranteed betting losses suffered by those who violate Conditionalization have no philosophical significance” (246). HEDDEN 2013a expands on this conclusion, arguing that we should treat distinct time slices like distinct agents in many cases where their collective action is against their collective interest. HEDDEN 2013b advocates “moving to an independently motivated picture of rationality which treats any *intrapersonal* requirements of rationality as deriving from more general requirements that apply equally in the *interpersonal* case” (36). Hedden goes on to advocate claims that resemble both definitions of time-slice epistemology, often defending them simultaneously.

Both definitions of time-slice epistemology have evolved as responses to the aforementioned challenges for traditional norms. It is not a coincidence that the theories are responsive to the same challenges. The first and second definitions are naturally related, since synchronic and impersonal norms are intimately connected. If the rational import of your past and future attitudes is mediated by your current opinions about those attitudes, then your opinions about the attitudes of other agents will often have that same import. Conversely, if the rational import of other agents is mediated by your opinions about their attitudes, then your opinions about your past and future attitudes will often have that same import.

However, it is important to recognize that while the first and second definitions of time-slice epistemology are connected, they are indeed independent claims. For starters, some norms are essentially synchronic and personal. The standard principle of Reflection defended by VAN FRAASSEN 1984 is one example. Reflection demands that your current credence in a proposition match your expected future credence in that proposition. The norm is synchronic, since your expectations of credences are among your current mental states. But Reflection is personal, as it assigns special rational import to your expectations of your own future credences, as opposed to the future credences of other agents. The Qualified Reflection norm defended by BRIGGS 2009 can helpfully direct us to impersonal revisions of Reflection. Qualified Reflection says roughly that if it is given that you will later have some particular credence in a proposition as a result of rationally updating on veridical evidence,

then you should already have that very credence in that proposition. This norm naturally follows from an impersonal norm. Suppose that you are certain that some agent with all your evidence has rationally updated on some additional veridical evidence. Then given that she has some particular credence in a proposition, you should have that very credence in that proposition. This norm applies whether the agent in question is someone else or some other time slice of yourself. In developing this impersonal replacement for Reflection, we are not advancing our first notion of time-slice epistemology, as each of the norms just considered is synchronic. But we are advancing the second notion of time-slice epistemology, and thereby addressing many of the same concerns that motivated our earlier rejection of diachronic norms.

In addition to synchronic personal norms, some theorists may accept impersonal norms that are essentially diachronic in the intrapersonal case. BURGE 1993 is one example. Burge argues that getting information from your past self is like getting information from other agents, as both memory and testimony involve “purely preservative” processes that directly transfer justification from one self to another. Suppose you have a justified belief and that I get this belief from you by testimony. Burge argues that my justification for my belief may have nothing to do with facts about me, such as the fact that I just heard you express the belief or the fact that I believe that you are reliable. The warrant may instead be just the same warrant that you have for your belief. In the same way, my inherited justification for remembered beliefs may have nothing to do with facts about my current mental states. Hence Burge may accept the second notion of time-slice epistemology without endorsing the first. The same goes for LACKEY 2008 when she compares memory and testimony, while arguing that “it is not enough for testimonial justification or warrant that a hearer have even epistemically excellent positive reasons for accepting a speaker’s testimony—the speaker must also do her part in the testimonial exchange by offering testimony that is reliable or otherwise truth-conducive” (155).

Both definitions of time-slice epistemology are natural and important. Neither has a stronger claim to fame. By contrast, it is worth considering a third definition that may initially seem attractive, namely the thesis that what is rational for your current time slice supervenes on its intrinsic properties. This definition is admirably simple, but ultimately less natural and compelling than those we have considered so far. For instance, the definition entails that time-slice epistemology is simply off-limits for many externalists, including anyone who accepts both that what you should believe depends on what knowledge you have and also that intrinsic duplicates can differ with respect to what knowledge they have. This is an unwelcome result, as the spirit of time-slice epistemology is intuitively independent of debates over epistemic

externalism. In addition, it is not clear that the intrinsic properties of time slices even include the sort of mental properties on which normative facts are meant to supervene. It may well be that strictly speaking, instantaneous temporal parts of objects have mental states only in virtue of having certain extrinsic properties. By contrast, the synchronic and impersonal definitions of time-slice epistemology merely require that there are facts about what mental states you are in at a particular time. These definitions allow that you may have a certain desire at a particular time partly in virtue of properties that you satisfy at other times, just as you may be painting a house at one instant in time partly in virtue of properties that you satisfy at nearby times. The same goes for your beliefs, memories, and other mental states.<sup>4</sup>

## 2 A case study: norms for action under indeterminacy

Having addressed the first programmatic goal of this paper, I will now turn to assessing particular time-slice theories. Our central case study begins with a thought experiment from VAN INWAGEN 1990:

Suppose that a person, Alpha, enters a certain infernal philosophical engine called the Cabinet. Suppose that a person later emerges from the Cabinet and we immediately name him 'Omega'. Is Alpha Omega?... Let us suppose the dials on the Cabinet have been set to provide its inmates with indeterminate adventures. (We need not agree on what would constitute an indeterminate adventure to suppose this. Let each philosopher fill in for himself the part of the story that tells how the dials are set.) Alpha has entered and Omega has left. It is, therefore, not definitely true or definitely false that Alpha is Omega. (243-4)

Following WILLIAMS 2013, we can use this story to raise questions about how agents ought to act in indeterminate decision situations.<sup>5</sup> Say that you are Alpha. You are walking toward the Cabinet. It is indeterminate whether you will survive what is about to happen inside it. How should you make decisions about the future? Say that a broker offers you an investment: if you pay him 10 dollars now, he will pay Omega 25 dollars when Omega comes out of the Cabinet.<sup>6</sup> Should you take the bet?

---

4. A referee worries that some norms govern "movements of mind," such as the norm requiring you to *change* your beliefs when they are inconsistent. But the point of time-slice epistemology is that we can derive such requirements from synchronic, impersonal norms. It is a fundamental fact that you are required to have consistent beliefs. This fact entails the less fundamental fact that you must reject some of your earlier beliefs when those beliefs were inconsistent, just as you must reject inconsistent beliefs held by others.

5. For readers skeptical about whether there are any genuinely indeterminate decision situations, the arguments of this paper may be more felicitously applied to situations where agents have imprecise credences about facts relevant to their decisions (cf. §3).

6. For simplicity, I assume throughout that agents are certain of the relevant details of their decision situations. In order to sidestep concerns about whether you could spend 10 dollars before entering the Cabinet, we could imagine the broker adding or subtracting from your immediate felt pleasure rather than from your bank account.

Standard decision theory does not answer the question for you. Williams argues that a complete account of indeterminacy should tell you how to respond to the broker, i.e. which if any responses are permissible, and which if any are obligatory.

The decision theory defended in WILLIAMS 2013 yields a straightforward verdict about the broker case, namely that it is okay for you to take the investment, and okay for you to reject it. In fact, when it comes to isolated, one-off decision situations, the decision theory is simple. In any case where the supervaluationist says that some indeterminate claim has multiple sharpenings, it is permissible for you to act as if any of those sharpenings is certainly correct. In other words:

(CAPRICE) An isolated action is currently permissible for you just in case there is some sharpening such that the action has highest expected utility according to your current utility function and your current conditional credence function, conditional on that sharpening being correct.

In the Cabinet case, the indeterminate claim relevant for your decision is the claim that you are identical with Omega. The claim has two sharpenings: either you are Omega, or you are not. The first sharpening sanctions your taking the investment. The second sanctions your rejecting it. Hence we may conclude that either action is permissible for you.<sup>7</sup>

Caprice is restricted to isolated decisions, cases where an indeterminate claim is relevant for your decision but has never before been relevant for any others. Hence Caprice is not a complete decision theory, because it does not yield verdicts about diachronic decision cases. Say that after you accept or reject the investment, the broker offers you a loan. He will immediately pay you 15 dollars, in exchange for charging Omega 25 dollars when Omega comes out of the Cabinet. The investment and the loan together constitute a great pair of bets. The payoffs are just the same as in the bets discussed in ELGA 2010. If you accept both the investment and the loan, you will end up 5 dollars ahead: you immediately pay 10 dollars and receive 15, while Omega pays and receives 25 dollars later. Caprice tells you that you may accept or reject the investment. But it does not say what you should do about the loan.

There are multiple ways of extending Caprice into a more complete decision theory. For instance, our complete theory may say that your previous decision does not constrain your current rational actions. To state the norm precisely, let us say that an action is currently *sanctioned by a sharpening* for you just in case it maximizes util-

---

7. There is a further complication in the theory defended in WILLIAMS 2013, namely that strictly speaking, your action must be randomly chosen from among the actions sanctioned by sharpenings. This complication does not affect my arguments, and so I will set it aside for sake of simplicity.

ity according to your current utility function and your current conditional credence function, conditional on that sharpening. Then we may extend Caprice as follows:

(LIBERAL) An action is permissible just in case it is sanctioned by some sharpening.

This norm is implicitly indexed: an action is permissible for an agent at a time, and a sharpening sanctions an action for an agent at a time. According to Liberal, the condition mentioned in Caprice is necessary and sufficient for the permissibility of any action.

By contrast, we could instead expand Caprice by saying that if you rejected the investment, you cannot also reject the loan. This extension of Caprice is inspired by the following three claims in WILLIAMS 2013:

[A]gents should strive to make their actions dynamically permissible. (16)

We call an action *dynamically permissible* at time  $t$  just in case it maximizes utility on some sharpening live at the score at  $t$ . (16)

When an action is carried out that is permissible on some but not all sharpenings, the score updates by eliminating those on which it is not permissible. (16)

Here is one interpretation of these passages: let us say that a sharpening is *live* for you just in case it sanctioned all of your past actions at the time at which you did them. Then we may extend Caprice as follows:

(RESTRICTIVE) An action is permissible just in case it is sanctioned by some live sharpening.

Like Liberal, Restrictive is implicitly indexed. In particular, a sharpening is live for an agent at a time. The idea is simple: if your previous actions were sanctioned by some sharpenings but not others, then an action is currently permissible for you just in case it is consistent with how you acted before. If you rejected the investment earlier, then also rejecting the loan is not sanctioned by a live sharpening, and so it is not dynamically permissible, and so you should not do it. According to Restrictive, what you did earlier constrains what actions are currently permissible for you.<sup>8</sup>

---

8. If you have already rejected the investment, some theorists may say that rejecting the loan is permissible, while the joint act of rejecting the investment and the loan is not (cf. Caprice in WEATHERSON 2008 and Sequence in ELGA 2010). If that is right, then intuitively agents should be interested in which actions are such that performing them will not make it the case that you have performed an impermissible sequence of actions, and readers may interpret 'permissible' in the text as denoting this property.

There are two natural ways of implementing Restrictive. The set of live sharpenings may be independent of your current mental states, in which case Restrictive will clearly conflict with our first notion of time-slice epistemology. By contrast, it may be that your current mental states are rationally constrained so that they determine which sharpenings are currently live for you. For instance, it may be that your opinions about indeterminate propositions are rationally constrained to evolve over time in a way that reflects how you have already acted on those propositions. In that case, Restrictive will be compatible with our first notion of time-slice epistemology, since which actions are permissible for you will supervene on facts about your current mental states. But that is only because we will have accepted another constraint incompatible with time-slice epistemology: which opinions are rationally permissible for you will not supervene on facts about your current mental states, but will be partly determined by independent facts about how you have acted before. Hence either way, fans of Restrictive will end up endorsing some norm at odds with our first notion of time-slice epistemology. The same goes for our second notion. Restrictive resembles Reflection. As you deliberate about what you should believe, Reflection demands that you assign a special normative status to certain future beliefs, simply because they are *your* future beliefs. In the same way, as you deliberate about how you should act, Restrictive assigns a special normative status to certain past actions, simply because they were *your* past actions.

To sum up the dialectic so far: both Liberal and Restrictive entail Caprice, while also saying something about how you may act in repeated decision situations. Liberal is a synchronic and impersonal norm. The set of sharpenings of some indeterminate claim need not have anything to do with your actions at other times, or indeed with you in particular. Hence according to Liberal, your rational options are independent of your earlier actions, just as they are independent of the actions of any other agent. In order to implement Restrictive, you must accept some norm at odds with time-slice epistemology. The choice between Liberal and Restrictive is an illuminating case study in the development of time-slice epistemology. I have already made some general remarks in favor of time-slice theories. In the next section, I will defend Liberal against specific objections often raised by advocates of norms like Restrictive.

### 3 Answering arguments against Liberal

Caprice, Liberal, and Restrictive are theories about how you should act when faced with indeterminacy. They have close cousins, namely theories about how you should act when your evidence is limited. For example, suppose that you are merely deeply

ignorant of what happens inside the Cabinet. There is a determinate fact of the matter about whether you will survive, but your evidence does not support a particular precise credence about your survival. Perhaps many precise credence distributions are rational given your evidence, or perhaps your evidence licenses only imprecise credal states that contain many precise credence distributions as members. Either way, your evidence does not uniquely determine how you should act. The analog of Liberal for imprecise agents says that an action is permissible just in case it is sanctioned by some member of your imprecise credal state, while the analog of Restrictive says that an action is permissible just in case it is sanctioned by some member that also sanctions all your previous actions. These theories have been discussed in recent literature: ELGA 2010 and WHITE 2010 present challenges for both, while WEATHERSON 2008 explores a third option that combines attractive features of each. This paper focuses on Caprice and Liberal, but many of my arguments apply equally to decision theories for imprecise agents. The critical comments in this section are largely inspired by literature on the latter.<sup>9</sup>

There are a couple of arguments against Liberal that have nothing to do with betting. These arguments are scarce in print but common in conversation, and so they merit some discussion here. The first argument is that it could not be rational to first act according to one sharpening and next according to another, without some *reason* for changing how you act. If you first act as if you will survive the Cabinet, for instance, then you cannot start to act otherwise for no good reason. This argument may seem compelling, until we observe that similar reasoning yields conclusions that contradict standard principles of decision theory. It is widely accepted that if multiple actions each have maximal expected utility, then each of the actions is permissible, regardless of whether you have chosen between just these actions before. It is permissible to act one way and then another, with no reason for changing how you act. Since some groundless switching between alternatives is permissible, it cannot be that Liberal is incorrect merely in virtue of permitting some groundless switching.

In addition, it is not clear why the argument under consideration is any better than an analogous argument for the opposite conclusion, namely that you cannot rationally keep acting in the same way without some reason for *continuing* to act that way. The fact that you acted some way before has no intrinsic epistemic significance.

---

9. For all I have said, it may be that some cases of action under indeterminacy are best treated with theories that govern imprecise agents, while others deserve another treatment entirely. For instance, sometimes it may be *indeterminate* whether actions are permissible when those actions lead to indeterminate outcomes. DOUGHERTY 2013 argues that cases of vagueness generate indeterminate ethical judgments. RINARD 2013 defends an alternative decision theory according to which it is often indeterminate whether actions are permissible for imprecise agents. I regret that I cannot explore these theories in more detail here.

Insofar as you must have reasons why you are acting according to some sharpening, the mere fact that you have acted on some sharpening before constitutes just as much reason as the mere fact that you have never acted on some sharpening before. This point about epistemic reasons can be made more clear by comparing it with the same point about moral reasons. The fact that you have done some action before may be good evidence that the action is morally permissible, provided that you are a good person. But that fact itself does not usually have intrinsic moral significance in your deliberation about which actions are currently permissible for you.

The second common argument against Liberal is that the theory is impractical or even impossible to implement. The complaint is that Liberal requires that you reconsider every decision made in the face of indeterminacy, constantly reassessing actions from moment to moment. It is plausible that ordinary agents cannot constantly reassess their actions. If ought implies can, then we may conclude that Liberal is false. The same sort of objection applies to several other time-slice theories. It may not be possible for agents to reinvent their credences from moment to moment, for instance, as may seem required by the updating norms defended by WILLIAMSON 2000 and MOSS 2012.

To respond: this objection misidentifies the subject matter of Liberal and other time-slice theories. The norms articulated in an epistemology classroom govern deliberating agents. Time slices that are not deliberating are simply not in the scope of Liberal. It may be true that people often chug along without deliberating, responding to any indeterminate claim as they did before, without reconsidering what sharpening they are acting on. It may even be true that people cannot survive without acting in this way. But this does not challenge norms that tell agents what they should do when they do deliberate. To compare: it may be true that people often fall asleep and hence fail to consider or assess any reasons at all, and it may even be true that people cannot survive without sleeping. But this fact about human nature does not challenge ordinary norms governing lucid agents.

The most compelling argument against Liberal is pragmatic. Recall that Liberal says that you can reject each of a pair of bets when accepting both would guarantee you sure money, and when you do not gain or lose any evidence about the bets as they are offered to you. In a recent paper about agents with imprecise credences, ELGA 2010 claims that results like these are unacceptable. Elga explains: “rejecting both bets is worse for you, no matter what, than accepting both bets. And you can see that in advance. So no matter what you think about [whether you will survive the Cabinet], it doesn’t make sense to reject both bets” (4).<sup>10</sup> This argument resembles

---

10. JOYCE 2011 develops a related argument against Liberal and endorses a consistency requirement that

a diachronic Dutch book argument. Liberal does not force agents to act on different sharpenings over time, and so agents following Liberal cannot be guaranteed to lose money. But it appears that rational agents should not only prefer keeping money over losing it, but also prefer gaining money over gaining nothing at all.

This pragmatic argument depends on the claim that it is impermissible for informed agents to forego sure money. But on reflection, this claim is not so clear. In fact, sometimes it seems that we ought to be forgiving of agents who forego sure gains. The most familiar examples of this phenomenon do not involve agents who are torn between beliefs, but agents who are torn between values.<sup>11</sup> To take an example from SARTRE 1946: you must either join the Free French as a soldier in England, or stay home to care for your ailing mother. Suppose that after several days of agonizing reflection, you board a train for England. But on the train, you have a change of heart. The situation has not changed, i.e. you have just the same evidence and just the same values as you did the day before. But you regret joining the army, and you feel resolved to care for your mother. In this situation, it seems perfectly permissible for you to get off the train to England and head home. In the literature on moral dilemmas, several authors aim to predict this sort of result. RAZ 1997 argues that when facing moral dilemmas, “we are within our rights to change our minds” (119). BROOME 2000 argues that you may “make the best of a bad job” in such situations (34). It is true that if you return home, you could have done better overall. Instead of buying train tickets, you could have saved your money and definitely come out ahead. But that does not mean that you are inextricably bound to your decision from the moment you board any train.

The same goes for situations where you are deeply torn between options, not because you are torn between values, but because you are torn between beliefs. For example, suppose that your friend is vacationing out of the country, and that some horrible wildfires have started to destroy the city where he lives. There are several family photograph albums in his apartment and several valuable manuscripts in his university office. It is impossible for you to save both, and impossible for you to contact your friend to ask which he would prefer you to save. Suppose that after several minutes of agonizing reflection, you board a subway train for his office. But on the train, you change your mind. The situation has not changed, i.e. you have just the same evidence about your friend as you did before. But you regret heading for the manuscripts, and you have made up your mind that your friend would rather

---

“ensures that at least one of Elga’s bets will always be chosen by any agent using any reasonable decision rule” (316).

11. See MOSS 2014 for a more detailed development of this response to the arguments in ELGA 2010. For further connections between imprecise agents and agents with incommensurable values, see the discussion of insensitivity to evidential sweetening in SCHOENFIELD 2012.

you save the photographs. In this situation, it seems perfectly permissible for you to head for his apartment. It is intuitively permissible even if it means that you end up eating the cost of your subway ticket. The ticket is a sunk cost, and you may rationally ignore it. If this intuitive judgment is right, it is not always impermissible for informed agents to forego sure money.<sup>12</sup>

To sum up so far: there are many ways to be deeply torn between options. Agents may be torn between values or between beliefs, and they may be torn between beliefs because they have limited evidence or because they recognize that there is no fact of the matter about some question. It is not clear that agents in these situations are strictly forbidden from changing their minds. In fact, we are intuitively disposed to forgive some agents who forego sure money, even when their change of heart is not prompted by any change in their evidence. I do not mean to suggest that the moral dilemmas literature univocally supports our intuitions about the above cases. For instance, CHANG 1997b argues that rational agents cannot have incommensurable values, precisely on the grounds that practical reason prohibits agents from being “merit pumps” (11). But it is telling that Chang focuses on an example in which an agent accepts unfortunate trades with no hesitation or reflection. In general, we are most inclined to reject apparent mind changing as irrational when it happens quickly, unreflectively, repeatedly, or for strategic reasons. These intuitions can be comfortably accommodated by a theory according to which changing your mind is not itself impermissible, namely because the salient features of these cases may provide evidence that they do not involve the same sort of genuine changes of mind exhibited by agents in the Sartre case and the wildfire case. By contrast, it is more difficult for blanket injunctions against mind changing to accommodate the intuition that changing your mind can sometimes be okay.<sup>13</sup>

From a third-person perspective, we sometimes forgive agents for changing their minds. But there is one special sense in which mind changing never seems permissible. From a situated first-person perspective, changing your mind may seem wrong for first-order reasons. For instance, it may seem that you should not change your opinion about a proposition because you would then have the wrong opinion about

---

12. To clarify the dialectic: some theorists accept that your evidence uniquely determines which precise credence function you should have. Fans of this uniqueness claim have a ready response to pragmatic arguments against Liberal, namely that rational agents are never in situations where Liberal recommends foregoing sure money. The point of the present discussion is that Liberal and similar time-slice theories can and should be accepted even by those who reject the uniqueness claim.

13. The arguments here and in Moss 2014 are limited: rather than arguing that mind changing is always permissible, I defend mind changing against norms that say that it is never permissible. This defense is incompatible with blanket injunctions against mind changing, but compatible with more nuanced theories according to which rational agents can have mental states that preclude mind changing. For further discussion of such theories, see HINCHMAN 2003, KORSGAARD 2008, HOLTON 2009, BRATMAN 2012, and FERRERO 2012.

that proposition. The foregoing permissive theory of mind changing can accommodate this sort of judgment. If you are acting according to some sharpening, other actions will always seem wrong simply insofar as they are not sanctioned by that sharpening. If you are acting as if you will survive the Cabinet, then accepting the loan will seem wrong simply insofar as it does not have maximal expected utility for you conditional on your surviving. From a reflective third-person perspective, you may acknowledge that accepting the loan is sanctioned by some sharpening, which constitutes an important sense in which accepting the loan is just as permissible as rejecting it. But even as you reflect, you may act according to a sharpening that does not treat these actions as equally permissible. Here it is especially useful to think of different time slices as being like different agents. Say that you first reject the investment offered by the broker, and later reject the loan. Then your later slice will judge that your earlier slice acted incorrectly, and your earlier slice would have said the same about your later slice. But there is an important sense in which you may reflectively judge that this disagreement is faultless. This judgment is captured by our reflective endorsement of Liberal.

The notion of faultless disagreement is familiar from literature about agents with different tastes and values. For instance, we may simultaneously endorse something as beautiful or fun or valuable, while reflectively judging that other evaluations of it are not wrong. Here again, familiar literature about valuing can help us better understand what to say about acting under indeterminacy.

This connection is not an accident, but part of a larger pattern. It is reasonable to expect similarities between agents acting on incommensurable values and agents acting on indeterminate or imprecise credences. For one thing, some imprecise credences may themselves be the product of incommensurable values, namely incommensurable epistemic values. For example, it may be that different hypotheses about the Cabinet are supported by different prior probability distributions, where these priors encode incommensurable ways of balancing epistemic values such as strength, elegance, and simplicity. To make matters worse, it may ultimately prove difficult even to distinguish between the state of having incommensurable values and the state of believing that there is no precise fact of the matter about what is valuable. The mental states of believing and valuing may not function as independently as classical decision theory would have us believe, which could give us further reason to expect literature on incommensurable values to provide us with fruitful analogies for theories of action under indeterminacy.<sup>14</sup>

---

14. For further discussion, see LEWIS 1988 and PRICE 1989.

#### 4 Developing arguments against Caprice

Liberal is a promising theory of action under indeterminacy. But it is not perfect. Recall that Liberal entails Caprice, namely that an action is permissible in an isolated decision situation just in case it is sanctioned by some sharpening. WILLIAMS 2013 credits Elga with raising a problem for a close cousin of this latter principle.<sup>15</sup> Suppose that the broker offers you the same investment as before, only now he also offers you a third option. Instead of immediately accepting or rejecting the investment, you may choose to delay your decision. If you delay, the broker will offer you the investment again in five minutes, and pay you one dollar for waiting. Elga claims that in this situation, it is intuitively permissible for you to delay your decision. WILLIAMS 2013 agrees that delaying is intuitively permissible: “After all, Alpha won’t close off any of the rival options, and he’ll gain a dollar whichever way he goes” (19). The problem is that Caprice entails that when you have some significant credence that you would act on a different sharpening later, it can be impermissible for you to delay your decision. In particular, any sharpening that you could act on will fail to sanction delaying whenever the small amount you would gain by waiting fails to outweigh the expected possible loss of your making the wrong decision about the investment later.

In response, Williams says that this problem for Caprice arises only if we neglect some standard assumptions about the rationality of agents. He says that when assessing the permissibility of actions, we standardly assume that agents are certain that they are rational and will remain rational. From this assumption, Williams concludes: “the credences induced by the sharpening [that entails that Alpha is Omega] will say that rationality requires investing; and hence (given that they assume that the agent will do what is rational) those credences will assign full probability to the agent investing” (25). In other words, the part of you that believes that you will survive the Cabinet also believes that you will accept the investment later if you delay, and the part of you that believes that you will not survive also believes that you will reject the investment if you delay. If we understand Caprice as saying that you may act according to any of these opinions, then it is permissible for you to delay your decision. Williams is cheered by this result, and ultimately this argument constitutes the response to Elga that he most prefers.

Unfortunately, this interpretation of standard rationality assumptions yields several unhappy consequences. For starters, the interpretation entails that you are rationally *compelled* to delay your decision, since delaying always has highest expected

---

15. Elga actually raises a problem for the combination of a claim like Caprice and the claim that your action must be randomly chosen from among the actions sanctioned by sharpenings. Since I am interested in assessing Liberal, I will present a version of the problem that constitutes a challenge to Caprice itself.

utility according to your conditional credence function, conditional on some sharpening together with the claim that you will later act according to that sharpening. But intuitively, it is sometimes permissible for you to just go ahead and accept the investment from the broker. Insofar as part of you can say, "I should accept this investment, namely because I will probably survive to collect on it," that part of you can also say, "I should accept this investment as soon as possible, namely because I will probably survive and there is a real chance that I will miss out on a great investment if I delay." This intuition becomes stronger as delaying is accompanied by smaller sure gains and larger possible losses.

The simplest response to this objection would be to weaken the assumption that you must act as if you are certain of some sharpening and also *certain* that you will always act according to that sharpening. Perhaps we should assume only that you must act as if you are certain of some sharpening and also have at least some threshold credence that you will later act according to it.<sup>16</sup> Here is a precise proposal that could replace Caprice: consider the constraint that it is both certain that some particular sharpening is correct and also fairly likely that you will later act according to that sharpening. We could say that an action is permissible for you just in case it maximizes utility according to your utility function together with your credence function after it has been updated on some constraint of just this sort.

This tempered proposal would allow you to accept the investment from the broker. But the proposal shares some other unhappy consequences with the proposal that Williams defends. For example: intuitively, your decision about delaying could be informed by independent evidence about what would happen if you did delay. Let us suppose that if you were forced to accept or reject the investment, you would accept it. And suppose you know that when it comes to indeterminate questions about your survival, you tend to change your mind a lot. It is rare for you to form intentions and carry them out without vacillating. Then your self-aware opinions may recommend that you accept the investment immediately, while the artificial conditional credence functions relevant for the tempered proposal may recommend that you delay. In more generality: when faced with an indeterminate claim, you have an imaginary "mental committee" of sharp opinions about how you should act. The proposals considered so far mandate wishful thinking on the part of each of your mental committee members, e.g. each member is confident that you will do just the right thing if you delay your decision, accepting the investment if and only if you will survive the Cabinet. But that means your decisions may not be appropriately responsive to your evidence. Proposals that require you to act on sufficiently optimistic credence functions blunt

---

16. Robbie Williams suggested this response to me in personal communication.

the force of relevant information, such as facts about the likelihood and relative cost of your changing your mind.

The possibility of side bets raises a related problem. In the initial simple broker case, the proposal in WILLIAMS 2013 requires you to bet at any odds that you will either accept the investment and survive the Cabinet, or reject the investment and fail to survive. This requirement is counterintuitive; rationality should not mandate hubristic certainty about your decisions. It seems fine for you to hedge your bets, accepting significant side bets that pay off just in case you are making the wrong decision about investing. The permissibility of hedging is forcefully illustrated by other hypothetical decision situations. Suppose that instead of entering the Cabinet yourself, you are about to send your pet hamster Fluffy into the Cabinet. Fluffy is valuable to you, and ordinarily you would pay up to fifty dollars to ensure her survival. As Fluffy enters the Cabinet, you see that there is a lethal device attached to the exit door that will kill any creature that emerges. If it only costs fifty cents to disarm the device, it seems intuitively permissible for you to pay the fifty cents to save the creature that will emerge from the Cabinet. But if you are willing to pay the fifty cents, you are not intuitively also obligated to pay up to fifty dollars to disarm the device. As the price of disarming the device dramatically increases, you may eventually decide that it is not worth the money.

Another decision situation involving hedging highlights a final problem for Caprice. The situation comes from WEATHERSON 2008:

An agent is told (reliably) that there are red and black marbles in a box in front of them, and a marble is to be drawn from the box. They are given the choice between three bets.  $\alpha$  pays \$1 if a red marble is drawn, nothing otherwise,  $\beta$  pays a certain 45 cents, and  $\gamma$  pays \$1 if a black marble is drawn. (12)

Let us suppose that which color marble is drawn is fixed by some indeterminate claim, like the claim that you survive the Cabinet. Not everyone has clear intuitions about this decision situation. But some report that it does not seem irrational to prefer  $\beta$  over the other gambles. If that is correct, Caprice is incorrect. If you act on one sharpening, you act as if you are certain that a red marble was drawn, and so you prefer  $\alpha$  over the other gambles. If you act on the other, you prefer  $\gamma$ . Hence accepting the certain 45 cents is not sanctioned by any sharpening. If choosing  $\beta$  is rationally permissible, then we have yet another case where Caprice seems too uncompromising. WILLIAMS 2013 states that “action under indeterminacy does not tolerate compromise” (29). But our intuitions suggest that it should.

## 5 Alternative theories of action under indeterminacy

In light of these problems for Caprice, it makes sense to look for alternative theories about making isolated decisions in the face of indeterminacy. Then we can extend these theories to get alternative time-slice theories of action under indeterminacy, improvements on the Liberal theory defended in §3 above.

Here is one alternative to Caprice: instead of identifying your mental committee members with particular sharpenings, we could identify them with nuanced opinions about sharpenings. This proposal would unify what we say about cases where your evidence is indeterminate and about other cases where your evidence fails to uniquely determine how you should act. The idea is that in the former cases, you should act as you do when you have imprecise credences about determinate propositions. This imprecise credal state is a set of credence distributions, each of which assigns some precise credence to each sharpening of the indeterminate proposition under discussion. The accompanying decision theory resembles Caprice, only with members of the imprecise credal state standing in for sharpenings. In more detail: an isolated action is currently permissible for you just in case it maximizes utility according to your current utility function and some member of the imprecise credal state. For example, in addition to acting as if you will certainly survive the Cabinet or acting as if you will certainly not survive, you may act as if you think those outcomes are equally likely, as long as the relevant imprecise credal state contains this third opinion. The corresponding analog of Liberal is straightforward: in any decision situation, an action is permissible just in case it maximizes utility according to your current utility function and some member of the imprecise credal state.

This alternative decision theory is neutral about the nature of the imprecise credal state that determines which actions are permissible for you. There may well be additional bridge principles that constrain this imprecise credal state in light of your opinions about the relevant indeterminate propositions. For example, it may well be that if you are certain that some proposition is indeterminate, then you should act as if your imprecise credal state contains every single precise credence distribution over sharpenings of that proposition. On the other hand, we may sometimes feel compelled to say things like “it is determinate that you might survive the Cabinet, and determinate that you might not survive it, even though it is not determinate whether you will survive,” in which case we may feel justified in eliminating only very decisive precise opinions from the relevant imprecise credal state.<sup>17</sup>

---

17. In light of such bridge principles, one might worry about whether we are failing to distinguish action under indeterminacy from action on insufficient evidence. WILLIAMS 2013 rejects some theories of the former on the grounds that they do not assign a distinctive cognitive role to uncertainty induced by

The revised Liberal theory under discussion accommodates our intuitive verdicts about almost all of the §4 examples. For instance, some members of your imprecise credal state may be on the fence with respect to whether you should accept or reject the investment offered by the broker. Those members will prefer delaying the investment decision over making it, since you definitely gain something by waiting, and there is no cost associated with your making either particular decision later. Hence the revised Liberal theory allows that delaying your decision is sometimes permissible. In addition, the theory allows that simply accepting the investment is sometimes permissible, namely whenever your imprecise credal state contains members that are sufficiently confident that you will survive the Cabinet. Finally, it is permissible for you to merely pay fifty cents in the hamster case whenever your imprecise credal state contains members that have some credence that Fluffy will survive the Cabinet, but not enough credence to justify paying fifty dollars to save the creature that emerges. In fact, many members of your imprecise credal state will normally have this feature, which may be partly responsible for our intuition that being willing to pay fifty cents but not fifty dollars is an eminently reasonable disposition.

The revised Liberal theory calls for just one more point of clarification. Suppose that WEATHERSON 2008 is right that you may prefer getting a certain 45 cents over getting one dollar just in case you survive the Cabinet, and over getting one dollar just in case you do not survive. This preference is not sanctioned by any precise credence about your survival, according to standard decision theory. The problem is two-fold: our model of your mental state is not fine-grained enough, and standard decision theory hastily condemns certain sorts of risk aversion as irrational. Instead of identifying members of your credal state with precise credences, we should identify them with pairs of precise credences and *subjective risk functions*, measurements of risk aversion defined as in BUCHAK 2013. Instead of sanctioning just those actions that maximize expected utility, members of your credal state should sanction actions that maximize *risk-weighted expected utility*.<sup>18</sup> For example, risk-averse agents may have exactly .5 credence that they will survive the Cabinet, and yet rationally prefer getting a certain 45 cents over getting one dollar just in case they survive. Augmented in this way, the revised Liberal theory can accommodate our intuition that you may prefer the certain 45 cents, namely since such hedging will be rationally permissible as long as your credal state contains some sufficiently risk-averse members with middling

---

indeterminacy. However, it is not clear that we should expect norms of rationality to distinguish action under indeterminacy from all other sorts of action. There are multiple reasons why your evidence could fail to determine the likelihood of some outcome. It may be that the likelihood relation is not well-defined, or that the outcome itself is not well-defined. The failure of your evidence may matter for the purposes of evaluating your action, while the source of that failure does not matter at all.

18. BUCHAK 2013 develops and defends this permissive alternative to standard expected utility theory.

credences about your survival.

There are many further respects in which Liberal may be revised and expanded. For example, we could use the mental committee model to represent your *values* in addition to your credences. For instance, we could identify members of your mental state with combinations of precise credences, subjective risk functions, and value functions. Then your having incommensurable values might be represented by members of your mental state having distinct value functions. Conditional values might be represented by dependencies between the credences and values of your mental committee members.

In addition, we could endorse more general procedures for deriving normative facts from features of your mental state. For instance, it may be that your permissible actions are not restricted to actions sanctioned by some member of your mental state, but instead include actions sanctioned by reasonable aggregations of the preferences of those members. For example, suppose that some member of your mental state is certain that you will survive the Cabinet, and some member is certain that you will not survive. Then it may be permissible for you to act as if you are not confident of either claim, not because your mental state contains some third member with this moderate opinion, but because moderate actions are preferred by some reasonable aggregation of your immoderate preferences. The same goes for accepting the certain 45 cents. That action may be permissible, not because your mental state contains some risk-averse members, but because it is preferred by some reasonable aggregation of your immoderate preferences.

The theories I have defended are inspired by extant theories of agents with incommensurable values. If an agent is torn between stringent values, then intuitively she may sometimes act according to a moderate compromise of those values. That may be because she takes some moderate value function to be an additional legitimate expression of her character, or it may be that someone who identifies only with stringent values may nevertheless act according to a reasonable aggregation of those values. Either way, we can say just the same thing about agents acting in the face of indeterminacy. The real world is full of compromises. You may join the army while still visiting your mother every weekend. You may accept an investment from one broker while hedging your bets with another. When it comes to agents with incommensurable values, we have both intuitive and highly theorized judgments regarding mind changing, faultless disagreement, hedging, and compromise. These judgments provide us with fruitful resources for defending time-slice theories of action under indeterminacy.

## References

- ARNTZENIUS, FRANK. 2003. "Some Problems for Conditionalization and Reflection." *Journal of Philosophy*, vol. 100 (7): 356–70.
- BAYES, THOMAS. 1763. "An Essay Toward Solving a Problem in the Doctrine of Chances." *Philosophical Transactions of the Royal Society of London*, vol. 53: 370–418.
- BRATMAN, MICHAEL. 2012. "Time, Rationality, and Self-Governance." *Philosophical Issues*, vol. 22 (1): 73–88.
- BRIGGS, RACHAEL. 2009. "Distorted Reflection." *Philosophical Review*, vol. 118 (1): 59–85.
- BROOME, JOHN. 2000. "Incommensurable Values." In *Well-Being and Morality: Essays in Honour of James Griffin*, ROGER CRISP, editor, 21–38. Clarendon, Oxford.
- BUCHAK, LARA. 2013. *Risk and Rationality*. Oxford University Press, Oxford.
- BURGE, TYLER. 1993. "Content Preservation." *Philosophical Review*, vol. 102 (4): 457–88.
- CHANG, RUTH. 1997a. *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press, Cambridge.
- . 1997b. "Introduction." In CHANG (1997a), 1–34.
- CHRISTENSEN, DAVID. 1991. "Clever Bookies and Coherent Beliefs." *Philosophical Review*, vol. 100 (2): 229–47.
- DOUGHERTY, TOM. 2013. "Vague Value." Ms., Department of Philosophy, University of Sydney. Forthcoming in *Philosophy and Phenomenological Research*.
- ELGA, ADAM. 2010. "Subjective Probabilities Should Be Sharp." *Philosophers' Imprint*, vol. 10 (5).
- FERRERO, LUCA. 2012. "Diachronic Constraints of Practical Rationality." *Philosophical Issues*, vol. 22 (1): 144–64.
- VAN FRAASSEN, BAS C. 1984. "Belief and the Will." *Journal of Philosophy*, vol. 81 (5): 235–56.
- HEDDEN, BRIAN. 2013a. "Options and Diachronic Tragedy." Ms., Department of Philosophy, Oxford University. Forthcoming in *Philosophy and Phenomenological Research*.

- . 2013b. "Time-Slice Rationality." Ms., Dept. of Philosophy, Oxford University. Forthcoming in *Mind*.
- HINCHMAN, EDWARD. 2003. "Trust and Diachronic Agency." *Noûs*, vol. 37 (1): 25–51.
- HOLTON, RICHARD. 2009. *Willing, Wanting, Waiting*. Oxford University Press, Oxford.
- JOYCE, JAMES. 2011. "A Defense of Imprecise Credences in Inference and Decision Making." In *Oxford Studies in Epistemology*, TAMAR SZABÓ GENDLER & JOHN HAWTHORNE, editors, vol. 4, 281–323. Oxford University Press, Oxford.
- KORSGAARD, CHRISTINE. 2008. *The Constitution of Agency*. Oxford University Press, Oxford.
- LACKEY, JENNIFER. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press, Oxford.
- LEWIS, DAVID K. 1988. "Desire as Belief." *Mind*, vol. 97 (418): 323–32.
- MOSS, SARAH. 2012. "Updating as Communication." *Philosophy and Phenomenological Research*, vol. 85 (2): 225–48.
- . 2014. "Credal Dilemmas." Ms., Department of Philosophy, University of Michigan.
- PRICE, HUW. 1989. "Defending Desire-as-Belief." *Mind*, vol. 98 (389): 119–27.
- RAZ, JOSEPH. 1997. "Incommensurability and Agency." In CHANG (1997a), 110–28.
- RINARD, SUSANNA. 2013. "A Decision Theory for Imprecise Credences." Ms., Department of Philosophy, University of Missouri, Kansas City.
- SARTRE, JEAN-PAUL. 1946. "The Humanism of Existentialism." In *Jean-Paul Sartre: Essays in Existentialism*, WADE BASKIN, editor, 31–62. Carol Publishing Group, New York.
- SCHERVISH, M. J., T. SEIDENFELD & J. B. KADANE. 2004. "Stopping to Reflect." *Journal of Philosophy*, vol. 101 (6): 315–22.
- SCHOENFIELD, MIRIAM. 2012. "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences." *Philosophical Studies*, vol. 158 (2): 197–219.
- TALBOTT, WILLIAM. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies*, vol. 62: 135–50.
- TITELBAUM, MICHAEL. 2013. *Quitting Certainties: A Bayesian Framework for Modeling Degrees of Belief*. Oxford University Press, Oxford.

- VAN INWAGEN, PETER. 1990. *Material Beings*. Cornell University Press, Ithaca.
- WEATHERSON, BRIAN. 2008. "Decision Making with Imprecise Probabilities." Ms., Dept. of Philosophy, University of Michigan.
- WHITE, ROGER. 2010. "Evidential Symmetry and Mushy Credence." In *Oxford Studies in Epistemology*, TAMAR SZABÓ GENDLER & JOHN HAWTHORNE, editors, vol. 3, 161–86. Oxford University Press, Oxford.
- WILLIAMS, ROBBIE. 2013. "Decision Making Under Indeterminacy." Forthcoming in *Philosophers' Imprint*.
- WILLIAMSON, TIMOTHY. 2000. *Knowledge and its Limits*. Oxford University Press, Oxford.