

A Plea for Falsehoods

Juan Comesaña
University of Arizona

1. Introduction

Falsehoods can play two important roles: we can rationally believe them, and they can be reasons for us to form further beliefs or act. On a natural understanding of the relation between rational belief and rational action (an understanding to be defended here), the second role of false beliefs follows from the first. If it can be rational for us to believe falsehoods, then these falsehoods can rationalize further doxastic attitudes and actions. That falsehoods can be rationally believed is the received view in epistemology. Indeed, that falsehoods can be rationally believed was, until very recently, seen as an adequacy condition on epistemological views. If, for instance, a view had the consequence that a pre-Einsteinian belief in the additivity of speed is irrational, then that view was seen as inadequate in virtue of having that consequence. That traditional view is under attack from the “knowledge first” camp. The attack can take two forms. First, there are those who, while conceding that there can be false rational beliefs, hold that there cannot be false beliefs that satisfy certain other central epistemological notion (like the notion of epistemic justification), and, moreover, that rationality itself is to be explained in terms of this more central notion.¹ An example of this kind of view will be one which identifies epistemic justification with knowledge itself, and then explains rationality counterfactually as the property had by those beliefs that someone who is usually justified would believe. There is a whiff of condescension in the concessions some of these views make: “Sure,” their supporters want to say, “you can say that false beliefs can be rational if you are tender-hearted, but the real normativity lies on a factive notion.” A more directly imperialist kind of attack leaves behind even the symbolic concession, and holds that false beliefs can be at best excused, but never have any positive normative status.² In this paper I will deal primarily with this latter view, in part because, following Cohen (forthcoming), I do not think that the distinction between epistemic justification and (theoretical) rationality can do the job that the other approach requires it to do. Many of my arguments, however, can be applied to

¹In different forms, this kind of view is defended by Williamson (2000), Sutton (2005), Sutton (2007), Littlejohn (forthcoming) and Lasonen-Aarnio (forthcoming). Goldman (1986) also distinguishes epistemic justification from rationality, but doesn’t argue that only true beliefs can be epistemically justified. Susanna Schellenberg has defended the view that there are two kinds of evidence in cases of basic perceptual justification, and that one of them, but not the other, can justify falsehoods—see, for instance, Schellenberg (2013).

²This is the view put forward in Williamson (forthcoming).

the view that although falsehoods can be rationally believed they cannot be justifiably believed. My main argument for the claim that falsehoods can be rationally believed is simple: rational action requires rational belief, and we can sometimes act rationally on the basis of false beliefs. The main work done in this paper is in the refinement and development of that simple argument.

2. A Starting Point

According to a coarse-grained Humean conception of practical rationality, the rationality of actions supervenes on beliefs and desires. Thus, for example, it may be that it was rational for you to go to the movies because you wanted to meet your friend, and you thought that your friend was there. I will complicate the example as I go along, and the complications will make it more philosophically interesting, but already at this level of description the example raises some philosophical issues. For there are other philosophers, Kantians, who disagree with Humeans even at this level. Whether it was rational for you to go to the movies, Kantians would say, is not determined by your beliefs and desires. If, for instance, your desire to meet your friend is itself irrational, then an action based on it will not be rational. Humeans admit that some desires, instrumental ones, may be assessed for rationality. For instance, your desire to buy a movie ticket may be counted as rational, because it is based on your desire to go to the movies and your belief that in order to go to the movies you need to buy a ticket. But, plausibly, some desires are ultimate ends of yours, not instrumentally based on any further desires. This is the locus of the dispute between Kantians and Humeans: Humeans take these ultimate ends of yours as given, and assess the rationality of instrumental desires and actions in their terms; Kantians think that even the ultimate ends are assessable as rational or not.

I do not plan to solve this dispute between Humeans and Kantians. Rather, I am interested in something that Humeans and Kantians, as I am defining them here, have in common. Both believe that the rationality of actions is to be determined by the interaction of two things: some ultimate ends (the rational ones for the Kantian, whichever ones you happen to have for the Humean) and some aspect of your doxastic life. In what follows, then, I concentrate on this common core.³ For the sake of ease of exposition, however (and also, why not admit it, because it more closely aligns with my sympathies), I will proceed as if Humeans are right.⁴

³In case it is not obvious enough, I hereby declare that I am using the labels “Humean” and “Kantian” with only the barest historical interest in mind. Real Kantians, for instance, might think that the rationality of an action is determined by the universalizability of the maxim on which the subject acts. But even this view can be “consequentialized” by giving lexical priority to outcomes which are the result of actions based on universalizable maxims. In any way, I am interested only in views of practical rationality that are consequentializable in this way. Relatedly, notice that Campbell Brown’s examples of moral theories that are not consequentializable (those which are agent neutral, admit the possibility of dilemmas, or violate dominance) are less plausible as theories of practical rationality (see Brown (2011)).

⁴Kantians face the following problem: what is it rational for someone with irrational ultimate ends to do? They should say what I say about the analogous question in section 4.

As I said in the introduction, my main argument for the rationality of false beliefs can now be stated very succinctly. Much of the rest of the paper will be devoted to developing and defending it. The argument is the following: It would have been rational for you to go to the movies even if your belief that your friend would be there had been false. But your action would not have been rational had your belief been irrational. Therefore, your belief that your friend would be at the movies could have been both false and rational. This is an instance of the following more general argument:⁵

1. Necessarily, if a subject S rationally ϕ -s on the basis of S 's belief that p , then it is rational for S to believe that p .
2. It is possible for a subject S to rationally ϕ on the basis of S 's belief that p even if p is false.

Therefore,

3. It is possible for there to be false rational beliefs.

There are a number of ways of resisting this argument. One can, for instance, deny premise 2, and hold that rational action does not tolerate false beliefs. I develop a precise version of this view in sections 5 through 8, and then criticize it in sections 9 through 11. The criticism in section 11 also applies to the more irenic view mentioned in the introduction, according to which it is true that there can be false rational beliefs, but false that there can be false justified belief. One can also deny premise 1, and in a couple of different ways. One way of denying premise 1 is to hold that rational action requires mere belief. I address this objection in section 4. More fundamentally, one can hold that the coarse-grained Humean picture of rational action that the argument relies on is fundamentally misguided and should be replaced by a fine-grained picture. I agree with this last complaint, but I do not think that it shows that there is something wrong with my argument. Even after replacing the coarse-grained Humean picture with a fine-grained one, it will still be true that rational action requires rational belief (and not just rational credences). Explaining why this is so will take some time. I begin the explanation in this section, and continue it in section 3. In section 12, I develop my view that rational action requires rational belief (and credences).

The objector is right that the coarse-grained nature of the Humean theory lands it in trouble. Suppose that you do desire to meet your friend at the movies, but you also desire to go to the park, and you can't do both at the same time. What now? If our picture of your pro-attitudes begins and ends with coarse-grained desires, then we could either say that you find yourself in a dilemma, because, whatever you do, you will frustrate one of your desires, or that you find yourself in a very happy situation (there's no good antonym for "dilemma"), because, whatever you do, you satisfy one of your desires. But rational choice is neither as hard nor as easy as that. In addition to (some would say "instead of") coarse-grained desires, we have fine-grained preferences. You prefer to meet your friend over not meeting him, and you also prefer going to the park

⁵The ϕ s in the premises are to be understood as variables (implicitly universally quantified) ranging over action-types—otherwise, the argument's second premise would be obviously question-begging.

over going to the movies. What should you do? Well, it depends: is your preference for meeting your friend stronger than your preference for going to the park? If so, go to the movies. If it's the other way around, go to the park. If the difference between the preferences is exactly the same, then you are indeed in one of those happy situations where you can do anything and be rational.

The Humean theory is coarse-grained not only because it deals in desires as opposed to graded preferences, but also because it deals in beliefs as opposed to graded credences. How do you know that your friend will be at the movies, and not at the park? You know, let us suppose, because he told you that he would go to the park if it doesn't rain, but go to the movies otherwise, and the weather report calls for 90% chance of rain. If our picture of your doxastic life begins and ends with coarse-grained beliefs, then you should go to the movies if you believe it will rain and go to the park otherwise. But in addition to (some would say "instead of") coarse-grained beliefs, you have fine-grained credences. Your degree of belief that it will rain is .9, let's say. Does this entail that you believe that it will rain? It's hard to say—we'll come back to that issue later. But your credences and preferences interact with each other in an intelligible way: the rational choice is determined by the strength of your preference for meeting your friend rather than not meeting him (as compared with the strength of your preference for the park over the movies) together with your credence about your friend's whereabouts. Your credence in your friend's whereabouts depends in turn on your credence in rain and your trust in what your friend told you.

Notice that there are two doxastic attitudes that are relevant to the rationality of your going to the movies: your credence in the proposition that it will rain, and your belief in the proposition that your friend will be at the park if it rains and at the movies otherwise. As we will see in the next few sections, credences and beliefs play two importantly different roles in standard decision theory (although the role that beliefs play is not generally discussed). Beliefs determine which states of the world you should take into account in planning your decision, and credences (together with preferences) determine which action is rational given those states.

Let us suppose that your preference for meeting your friend is high enough that, combined with your high enough credence that it will rain, you are rational to go to the movies. Let us also suppose that everything goes according to plan and you meet your friend. This is a standard case of decision under risk. You had two actions available to you: going to the park and going to the movies. You cared most about meeting your friend, and you knew that he would be at the park if it does not rain or at the movies otherwise. Moreover, you had very good information about whether it was going to rain. Given that this is a fair description of your choice situation, you did what any sane theory of rational choice would have told you to do. In this case, moreover, being rational payed off.

You cannot communicate with your friend now, and both the movies and the park are far away. This is why you have to rely on the weather report and your friend's testimony, instead of a more direct kind of coordination. Consider then a variation of the case where, against the odds, it does not rain, your friend goes to the park, and you do not meet him. It is agreed on all hands that although going to the movies was worse

for you than it would have been to go to the park, it was nevertheless rational. This should already be a reason to think that we can be rational in believing falsehoods. Arguably, truth (or perhaps knowledge) is to belief as bringing about the intended outcome is to action.⁶ Since we can act rationally without bringing about the intended outcome (for example, when something that was by our own lights unlikely happens), why couldn't we believe rationally without believing truly?

The argument for the rationality of believing falsehoods does not stop with that analogy. It continues (throughout this paper) with a second variation on our case. In this second variation, the weather reports were right and it did rain, but your friend unpredictably went to the park nevertheless. My starting point is that in this case, as in the previous one, the unfortunate consequences of your action do not impugn its rationality. In saying that this is my starting point, I do not mean to suggest that there are no arguments for it. Indeed, later on I will provide arguments for it. By calling it a starting point, rather, I mean to remind us that it is what most of us would believe before any theorizing. It may well be, of course, that theorizing should change pre-theoretical assumptions—that is what arguments are for. But it is nevertheless important to record what the pre-theoretical assumptions are.

My starting point is not just about this specific case and its variations. Rather, it is about structural features of them. The structural features are the following. A decision is to be made among a variety of possible actions. There is a relevant partition of the different ways the world might develop—the world is guaranteed to develop in exactly one of those ways. Let us call the combination of an action with a way the world might be a possible outcome of that action. The agent has opinions about the relative likelihood of each possible state of the world, and prefers certain outcomes to others (for instance, in our case you think it is more likely than not that it will rain, and you prefer meeting your friend to not meeting him). In every case, the agent acts in a way that makes sense in light of his opinions and preferences. In the base case, everything goes as planned. In variation 1, something happens that, by the lights of the agent, was unlikely to happen. In variation 2, the partition of states over which the agent's preferences are distributed was determined by false beliefs.

That is my argument that the claim that rational action requires rational belief does not depend on using a coarse-grained rather than a fine-grained theory of rational action, for even in a fine-grained theory there is still a place for belief. The argument so far was informal, but in the next section I show how it survives in a more formal environment.

⁶It would be more precise to say that truth (or perhaps knowledge) is to belief as maximizing utility is to action. Notice that I say "maximizing utility," not *expected* utility. Arguably, maximizing expected utility is to action as rationality is to belief. More on utility, expectation, and maximization in the next section.

3. The Role of Belief in Standard Decision Theory

According to what I will call “standard decision theory,”⁷ a decision problem is determined by a set of available actions, a partition of states of the world, a credence distribution C , and a utility function U over pairs of actions and pieces of the partition (which we can also call “outcomes”). It is assumed that the utility function is an *interval* scale in that the differences between preferences matter. For instance, if your utility for dulce de leche ice cream is 10, your utility for chocolate ice cream is 9, your utility for pasta bolognese is 8 and your utility for baked eggs is 6, then your preference for pasta bolognese over baked eggs is stronger than your preference for dulce de leche over chocolate ice cream. It is also assumed, however, that U is unique only up to positive linear transformations, so that the absolute numbers do not matter. Many economists and philosophers who use standard decision theory think that what the utility function measures is not some intrinsic “degree of desire” that you have over individual outcomes, but rather a preference relation over such outcomes. Representation theorems in decision theory guarantee that if an agent’s preferences over outcomes satisfies certain axioms (which are often assumed to be rationally required), then there is an interval utility function (unique up to positive linear transformation) with the following two properties: first, it “represents” your preferences in the sense that $U(a) > U(b)$ if and only if you prefer outcome a to b ; and second, the utility that it assigns to any “lottery” between outcomes is the expected utility of that lottery (more on what “expected utility” means below).⁸

The credence function C is assumed to be a probability function,⁹ and a partial two-place conditional probability function is definitionally introduced as follows:

Definition of conditional probability: $C(p|q) = \frac{C(p \wedge q)}{C(q)}$, whenever $C(q) > 0$, and undefined otherwise.¹⁰

⁷What I am here calling “standard decision theory” takes elements from both Savage-style and Jeffrey-style decision theory (see Savage (1954) and Jeffrey (1965)). It assumes independence (both causal and probabilistic) between actions and states (see also note 11), but does not otherwise treat them differently.

⁸See von Neumann and Morgenstern (1944). Other representation theorems also derive a probabilistic credence function from these preferences.

⁹This means that C is defined over a set of propositions which is closed under the Boolean operations (conjunction, disjunction and negation). Moreover, C satisfies the following axioms:

1. For any proposition p , $C(p) \geq 0$
2. For any tautology T , $C(T) = 1$
3. For any two incompatible propositions p and q , $C(p \vee q) = C(p) + C(q)$

The first two axioms have a book-keeping nature in that they establish that a probability function is normalized to 1 and with a lower bound of 0. All the interesting formal results will go through if we replace the 0 and the 1 with other numbers. The third axiom, however, is more substantial, for it requires probabilities to be finitely additive.

¹⁰Philosophers have worried that the standard definition of conditional probability as a ratio of unconditional probabilities hides substantive assumptions which might well be philosophically indefensible—for instance, you might well think that $C(p|p) = 1$, even when $C(p) = 0$ or is undefined. An alternative is to take conditional probability as a primitive and define the unconditional probability of a proposition as its probability conditional on a tautology. For more on this, see Hájek (2003).

We have, then, a utility function and a probability function. Using them, we can represent a decision problem in terms of a matrix:

Table 1: Decision matrix

	$C(S_1)$	$C(S_{\dots})$	$C(S_n)$
A_1	$U(A_1 \wedge S_1)$	$U(A_1 \wedge S_{\dots})$	$U(A_1 \wedge S_n)$
A_{\dots}	$U(A_{\dots} \wedge S_1)$	$U(A_{\dots} \wedge S_{\dots})$	$U(A_{\dots} \wedge S_n)$
A_n	$U(A_n \wedge S_1)$	$U(A_n \wedge S_{\dots})$	$U(A_n \wedge S_n)$

The row headings refer to the actions available to the agent and the column headings to the credences in the different pieces of the partition of possible states of the world relevant to the decision. We are assuming here that there are no probabilistic or causal correlations between the actions and the states, and so the credence is the same for a given state, no matter which action is assumed to occur in that state. If the states *do* depend on the actions, then we need to modify this part of the picture.¹¹

The utility function and the credence function take values over the same outcomes. Whenever we have a probability distribution as well as any other function defined over a partition, we can talk of the *expected value* of that other function. The expected value is the average of the values over the partition weighted by the probability assigned to each outcome. In the case of a decision problem, each act will have associated with it an *expected utility*.¹²

Standard decision theory is the theory according to which a rational action maximizes expected utility. Thus, decision theory is the contemporary descendant of a Humean theory of practical rationality as instrumental rationality, according to which the rationality of an action depends on whether it was caused by the right combination of beliefs and desires. Decision theory replaces desires by preferences and beliefs by credences, as we argued it should be done, and then defines rationality in terms of maximization of expected utility relative to those credences.

As I argued above, however, there is still a place for belief alongside credences in standard decision theory. The place for belief might be hidden in the way standard decision theory is often deployed, but it can be found if one knows where to look. The place to look is the shape that the matrix takes.¹³

¹¹ *Evidential* decision theory, as in Jeffrey (1965), replaces the unconditional credence that the state occurs by the conditional credence of the state given the action, whereas (some versions of) *causal* decision theory, as in Gibbard and Harper (1978), replaces it by the credence in the counterfactual *if the action had been performed, then the outcome would have resulted*.

¹² More generally, whenever we have a utility function and a credence function defined over a partition Π with pieces from Π_1 to Π_n , the expected utility of an action A will be given by the following formula:

$$\sum_{i=1}^n U(A \wedge \Pi_i) \times Cr(\Pi_i)$$

¹³ Weatherson (2012) makes the point that the shape of the decision matrix is determined by your full beliefs, but he thinks that the beliefs need to amount to knowledge in order to play that role. Weatherson, then, is a prime example of a philosopher who adopts what I later call, following Dutant (forthcoming), “knowledge-based decision theory.”

Let us take a look at how standard decision theory would model your choice to go to the movies. There are two relevant states of the world: the state where it rains and your friend goes to the movies and the state where it doesn't rain and your friend goes to the park. There are also two relevant actions you might perform: go to the movies or go to the park. Let us suppose that your credence that it will rain is .9 (and so your credence that it will not rain is .1). Let us also suppose that you prefer every outcome where you meet your friend to every outcome where you don't, but you prefer the outcome where you are by yourself at the park to the outcome where you are by yourself at the movies. We need to be more precise than that, however: we need to figure out the *differences* in your preferences. Let us suppose that although you do prefer meeting your friend to not meeting him no matter what, the difference between your preference for being at the park by yourself over being at the movies by yourself is much larger than the difference between your preference for being at the movies with your friend to being at the park alone. Remembering that the absolute numbers do not have any meaning beyond establishing this difference, let us assume that your preferences are as in the following matrix:¹⁴

Table 2:

	$Cr(\text{Rain} \wedge M) = .9$	$Cr(\text{No rain} \wedge P) = .1$
Park	8	9
Movies	9	1

In that case, the expected utility for you of going to the park is: $.9 \times 8 + .1 \times 9 = 8.1$, whereas the expected utility of going to the movies is $.9 \times 9 + .1 \times 1 = 8.2$, with the result that the rational choice is to go to the movies. Thus, although you much rather prefer to be alone at the park than alone at the movies, your credence that it will rain is high enough to rationalize starting a trip to the movies.

In building the matrix I said that there were two relevant states of the world: one where it rains and your friend goes to the movies and one where it doesn't rain and your friend goes to the park. But why aren't there two other relevant states of the world: one where it rains and your friend goes to the park, and one where it doesn't rain and your friend goes to the movies? The answer is that you believe what your friend told you: that he would go to the movies if it rains and to the park otherwise. What your friend told you is incompatible with the other two states, and so in believing him you do not count those states as relevant to your decision. You do not countenance a state where it rains and your friend goes to the park for the same reason that you do not countenance a state where it rains and you will be abducted by aliens on your way to the movies: you do not believe that either state will occur. Of course, you *might* mistrust your friend, or for some other reason not believe what he tells you. In that case, your decision might take into account all four states and not only two. Indeed, you might also believe that it is a live possibility that you will be abducted by aliens on

¹⁴The first column from the left represents the state where it rains and your friend goes to the movies, and the other column the state where it doesn't rain and your friend goes to the park, the rows represent the actions available to you (start a trip to the movies or to the park) and the cells represent the four possible outcomes.

your way to the movies, in which case you might take into account even more states. But you would never take into account *all* logically possible states. Some of those you do not believe will happen. The role of belief in traditional decision theory, then, is precisely to determine how coarse-grained the states that you should count as relevant to your decision need to be. Because you believe what your friend told you, the state where it rains need not be further fined-grained into two: one where your friend told the truth and he goes to the movies and one where he goes to the park anyway.

So, even though on the surface traditional decision theory is thoroughly fine-grained in that it deals with credences and preferences, there is still a place for belief as the state that determines exactly which decision problem you will be facing. C. I. Lewis was right after all: nothing is probable unless something is certain.¹⁵

Now, your beliefs thus determine which decision problem you are in fact faced with, but they do not answer the question which decision problem you *should* be faced with. If your beliefs are totally irrational (if, for instance, you have ample evidence that your friend is lying to you), then, even though your beliefs determine that you are facing a particular decision problem, that is not the problem you should be facing. The problem that you should be facing, I will argue, is determined not by the beliefs you actually have, but by the beliefs it would be rational for you to have.

As I said before, the position that rational belief is required for rational action is under pressure from two sides. On the one hand, some philosophers think that more than rational belief is required for rational action, the more prominent proposal here being that knowledge is required.¹⁶ On the other hand, some philosophers think that less than rational belief is required for rational action, the more prominent proposal here being that mere belief is sufficient. I start by arguing against this latter position.

4. Rational Action, Rational Belief, and Dilemmas

Parfit is an example of a philosopher who thinks that rational belief is not required for rational action, because mere belief is enough. He would say that for a belief of yours to rationalize an action it is enough if it would be a sufficient reason to act that way if it were true (see Parfit (2011)). Thus, in our example, even if you believe that your friend will be at the movies because the Ouija board said so, it is rational for you to go to the movies.

Why should we think this? After all, just as the default position in epistemology is that a belief can rationalize another belief only if it itself is rational (this is the principle that launches the venerable regress argument), why shouldn't the default position in

¹⁵Of course, the role that I claim belief is playing here can also be played by credence 1. Plausibly, though, we cannot just identify belief with credence 1, even if they coincide in these cases. Anyway, if someone wants to claim that there is no place for beliefs in decision theory, even though there is a place for maximal credence, I need not disagree. My arguments can be restated in terms of the rationality of assigning credence 1 to falsehoods.

¹⁶Or rather, some of these philosophers will say that more than so-called rational belief is required, if they identify rational belief with knowledge.

the theory of rational action be that a belief can rationalize an action only if it itself is rational? It is interesting to note, in this connection, that Parfit himself adheres to this principle for the case of beliefs: he holds that a belief cannot be (wholly) rational if it derives from an irrational belief. Why the difference between the practical and the theoretical case? Because, Parfit thinks, it is better not to say that actions (or desires) can be epistemically irrational. But someone who holds that rational action requires rational belief need not say that the actions based on irrational beliefs are *epistemically* irrational. I agree with Parfit that this wouldn't make much sense. They are just as practically irrational as actions that are based on rational beliefs but don't match them (like drinking what's on the glass when we want water and we rationally think that it contains gin). We don't think that beliefs that are based on irrational beliefs are irrational in a different sense from beliefs that based on rational beliefs but don't match them (like believing that Obama is in Canada based on the belief that he is in Paris)—why think that we would be forced to say that actions can be irrational in two different senses if we held the analogous view in the practical case?

At any rate, regardless of the internal consistency of Parfit's views, there is, I think, an argument for his view in the practical case that we should pay attention to (Parfit himself does not give this argument). The argument starts with a question that is supposed to be embarrassing for friends of the view that rational action requires rational belief. Applied to our case, the question is the following: if going to the movies is not the rational thing to do when you believe that your friend is there and you want to meet him (even when this belief is irrational), then what is the rational thing for you to do in that situation?

In response, we should start by clarifying that going to the movies might indeed be the rational thing for you to do, even if your belief is irrationally held. After all, it may well be rational for you to believe that your friend will be at the movies—if, for example, your friend told you as much. We need to distinguish between *ex-ante* and *ex-post* rationality, both for beliefs and actions.¹⁷ Different philosophers have different theories of *ex-ante* rationality for belief. But regardless of your view of the *ex-ante* rationality of belief, you need to distinguish it from the *ex-post* rationality of belief. A proposition might be *ex-ante* rational for you to believe regardless of whether you in fact believe it or not. But when you do believe a proposition, we can ask not only whether the content of that belief (the proposition believed) is *ex-ante* rational for you or not, but also whether the mental state itself of believing that proposition is rational. This *ex-post* rationality of belief depends not only on the *ex-ante* rationality of its content, but also on the basis for the belief itself. Thus, if both your friend and the Ouija board "said" that your friend will be at the movies but you believe only the Ouija board, then your belief that your friend will be at the movies is *ex-ante* but not *ex-post* rational.¹⁸ If we want, we can of course apply the same distinction to credences—your credence in a given proposition might be *ex-ante* but not *ex-post* rational, as will happen when (to put it roughly) you have good reasons for assigning it that credence, but you assign it for bad reasons.

¹⁷See Goldman (1979). His terminology is better than distinguishing between propositional and doxastic justification, for it can apply to actions as well.

¹⁸Pace Lehrer (1971).

All of that will be old news for epistemologists, of course, but an analogous distinction applies to the rationality of actions as well. Thus, an action-type is *ex-ante* rational just in case it maximizes expected utility relative to your *ex-ante* rational credences and beliefs, and it is *ex-post* rational just in case it maximizes expected utility relative to your *ex-post* rational credences and beliefs. “The rational thing to do” is ambiguous: it might refer to the *ex-ante* rational action or to the *ex-post* rational action. When your credences are *ex-ante* but not *ex-post* rational and you maximize expected utility relative to them, then your action is also *ex-ante* but not *ex-post* rational.

With this distinction in hand, let us now go back to the question that we asked on behalf of Parfit: if when you irrationally believe that your friend will be at the movies it is not rational for you to go to the movies, then what is it rational for you to do in that situation? In reply, we disambiguate. If we are asking which action is *ex-ante* rational, then it may well be that it is indeed to go to the movies. If we are asking which action is *ex-post* rational, then the answer is that none is. An action can be *ex-post* rational for a subject only if the subject’s beliefs and credences on which the action is based are themselves *ex-post* rational,¹⁹ and so when the subject’s credences are not *ex-post* rational, no action is *ex-post* rational either.

Is that an embarrassment for the claim that *ex-post* rational action requires *ex-post* rational belief? It seems to have as a consequence that there can be rational dilemmas. Say that an agent faces a rational dilemma just in case whatever he does his action is not rational. The view that *ex-post* rational action requires *ex-post* rational belief does have the consequence that agents may face a particular kind of rational dilemma. But not all dilemmas are equally repugnant, and the kind of dilemmas allowed for by the view that *ex-post* rational action requires *ex-post* rational belief has three properties which, together, make them acceptable.

First, as explained above, they are only *ex-post* dilemmas. There is, after all, an *ex-ante* action that is rational, which is to go to the movies.

Second, the dilemmas in question are *self-imposed*. At least since Aquinas, the distinction between self-imposed and world-imposed dilemmas has been thought very important. A dilemma is self-imposed when it is due to a previous action of the agent. Suppose, for instance, that you make two incompatible promises, and contrast that with the case where you make only one promise, and through no fault of your own it is now impossible for you to fulfill it. Some philosophers will of course claim that in neither case do you face a genuine dilemma. But many others will say that although you cannot face world-imposed moral dilemmas, it is perfectly possible for you to face self-imposed moral dilemmas. The kind of dilemma allowed for by the view that *ex-post* rational action requires *ex-post* rational belief is of the self-imposed kind. True, there need not have been an action (in the narrow sense of “action”) that you performed as a result of which you are in this dilemma, but there is indeed something you did (in a

¹⁹Plausibly, that the action maximizes expected utility relative to the subject’s *ex-post* rational credences and beliefs is only a necessary, not a sufficient condition for the *ex-post* rationality of the action. In addition, for example, we might want to require the action to be based in the right way on the subject’s credences and preferences. This basing requirement is typically absent in presentations of decision theory.

broad sense of “doing”), namely, not forming rational beliefs for the right reasons, as a result of which you now face a dilemma.

Third, and perhaps more importantly, the kind of dilemma in question is *recoverable*. You can, after all, form rational beliefs for the right reasons, in which case you will have *ex-post* rational beliefs, and in which case (in turn), there will be an *ex-post* action that is rational for you to perform: go to the movies.²⁰

To sum up our discussion so far: it is *prima facie* plausible to think that rational action requires rational belief. An objection to this view observes that, in cases where rational beliefs are absent, practical rational dilemmas ensue. But the dilemmas whose possibility the view in question allows are *ex-post*, self-imposed and recoverable. These kinds of dilemmas are not the problematic kind. In what follows I will be assuming, then, that rational action requires rational belief.

5. Subjective Bayesianism

But rational action requires more than rational belief. If all of your beliefs are rational, that takes care of the shape that your decision problem takes: it justifies you in using one particular matrix rather than another. But you may well have all rational beliefs but irrational credences. The *shape* of your problem might be rationally based, but its content might not. Suppose, for instance, that you rationally believe your friend, and so you rule out for consideration states where it rains but your friend goes to the park, or states where it doesn't rain but your friend goes to the movies nevertheless, but your assignment of a high credence to the proposition that it will rain is based on the results of the Ouija session. Just as with beliefs and actions, the *ex-ante/ex-post* rationality distinction applies to credences as well. In this case, your .9 credence that it will rain is *ex-ante*, but not *ex-post*, rational. For the same reasons that your action will not be *ex-post* rational unless your belief that your friend is at the movies is itself *ex-post* rational, your action cannot be *ex-post* rational if your credences are not *ex-post* rational. We need, then, some account of the rationality of credences.

If that is right, then standard decision theory is doubly incomplete, for it doesn't include an account of the rationality of either beliefs or credences. Subjective Bayesianism is a minimalist theory of the rationality of credences, and it includes an implicit and incomplete account of the rationality of beliefs, so I turn to it now.

According to subjective Bayesianism, there is a synchronic constraint and a comparative constraint on rational credences. The synchronic constraint was already alluded to, and it consists in the claim that, at any given time, the credences of a subject must obey the probability calculus. For many of us, this requires too much, for it deems as irrational (for example) our assigning less than maximal credence to any tautology, no matter how complicated. But my topic here is not logical omniscience, on which much

²⁰All of this is not to deny that the issue of “contrary to duty” conditionals (for instance, “if you believe that your friend is at the movies, then you ought to go to the movies”) is very interesting. For more on this, see my Comesaña (2016).

good work has been done.²¹ Closer to my topic is the opposite worry: that simply requiring probabilistic coherence as the only synchronic rational constraint on credences is too lenient, for it sanctions as rational patently irrational credences. More on this in the next section.

The subjective Bayesian comparative constraint has it that a subject's credence function at a time t' should be the conditionalization of his credence function at an earlier time t given all the new evidence that the subject has accumulated between t and t' :

$$\text{Conditionalization: } C_{t'}(p) = C_t(p|E)^{22}$$

According to subjective Bayesianism, that is *all* that is required for credences to be rational. No other constraints are imposed on credence functions. As long as they are probabilistically coherent and those adjacent in time are related to one another by conditionalization, the subject whose doxastic life they represent is rational.

Does subjective Bayesianism have anything to say about the rationality of beliefs? Not directly. But it does indirectly, at least if we are willing to hold that if a subject assigns credence 1 to a proposition, then he believes it (notice that assigning credence 1 is treated here as a sufficient, but not necessary, condition on belief). If so, given its conditionalization constraint, according to subjective Bayesianism all of the propositions that are part of a subject's evidence are rational for that subject. For, according to the conditionalization requirement, all of the propositions that are part of a rational subject's evidence receive credence 1, and according to the proposed sufficient condition on belief that means that all of the propositions that are part of a rational subject's evidence are believed by that subject.

So, to summarize: according to subjective Bayesianism all it takes for credences to be rational is for them to be probabilistically coherent and to evolve by conditionalization. Moreover, it is a sufficient condition for a belief to be rational that it be a part of the subject's evidence.

6. Objections to Subjective Bayesianism

Subjective Bayesianism (as well as any decision theory based on it) is subject to well-known objections. The obvious one is that it is too subjective: that the constraints it imposes on rational credences are too lenient. For instance, it is perfectly compatible with subjective Bayesianism that you have opinions which are representable by a credence function which, in our example, assigns a very high credence to the proposition that the moon is made of cheese. For this credence assignment to be compatible with

²¹See, for instance, Stalnaker (1991), Stalnaker (1999), Williamson (forthcoming), Christensen (2004).

²²Notice that I didn't call the conditionalization constraint "diachronic" but merely "comparative," and for good reasons. "Diachronic" suggests that the constraint imposes requirements on the *process* by which credence functions should evolve over time, but no such requirement is imposed by conditionalization. As far as it is concerned, it doesn't matter *why* your credence functions at two times obey it, only that they do.

subjective Bayesianism, that credence function must be the result of conditionalizing on whatever your evidence at that moment is. Now, subjective Bayesianism is also subject to the objection that it has no account of evidence (more on that below), but let us assume for the time being that your relevant evidence is that your friend will be at the movies if it rains (and at the park otherwise). The conditionalization requirement, then, has the consequence that your credence function before acquiring this evidence must have assigned a high credence to the moon's being made of cheese, conditional on the proposition regarding the whereabouts of your friend. This is, of course, a bizarre conditional credence to have, but that is precisely the point of the objection: according to subjective Bayesianism, such a credence function would be perfectly rational. Of course, as things stand, if you act on the basis of those credences you will reliably not get what you want. But, for the subjective Bayesian, this would just be a case of bad luck, where things do not go according to your expectations—whatever those are.²³

Another problem for subjective Bayesianism resides not in its subjectivity, but in the requirement of conditionalization. Notice, first, that any evidence that the subject conditionalizes on automatically receives credence 1 (the conditional probability of any proposition p given itself is 1). This may already be problematic for some philosophers, but it is not the main problem with conditionalization. Rather, the main problem with conditionalization is that whenever a proposition acquires credence 1, it retains credence 1 forever. This is so because the probability of any certain proposition on any other proposition (with positive credence) is 1. Thus, for the Bayesian, evidence is *cumulative* in the sense that once a proposition is part of a rational subject's evidence, it remains in that subject's evidence forever. The requirement of conditionalization thus conflates forgetting with irrationality. But not only is this requirement incompatible with forgetting, it is also incompatible with the widespread phenomenon of "knowing less by knowing more." Suppose, for instance, that we see that one black and one red ball are introduced in an urn.²⁴ Thus, it is now part of our evidence that at least one ball is not black. But then we perform the following experiment: we take out one ball, note its color and put it back in the urn. We then shake the urn and perform the same experiment again. Suppose that we do this 10,000 times, and that each time the ball is black. The rational response, in this case, is to doubt our previous thought and suspect that what looked like a red ball was actually another black ball. It would be irrational to insist on being certain that there is a red ball after such overwhelming evidence against that proposition. However, if we update our credences by condition-

²³Some subjective Bayesians are fond of appealing to "convergence theorems" to fend off the subjectivity objection. According to these theorems, subjects who update their credences by conditionalization and who receive the same evidence tend to converge on the same credences. The philosophical importance of these theorems, however, is greatly diminished by two considerations. First, they are "limit" theorems about what the credences tend to as the evidence grows unbounded. As such, the theorems are compatible with any two subjects who conditionalize on the same evidence having radically different credences given any finite amount of evidence. Second, and more to the point, the theorems are subject to the constraint that the respective credences must be in considerable agreement regarding the conditional probabilities of propositions on the evidence. That this must be so is clear in light of the following fundamental result: for any $n \in [0, 1]$ and credence function $C_{t'}$ such that $C_{t'}(p) = n$ and $C_{t'}(E) = 1$, there is a credence function C_t such that $C_{t'}$ results from conditionalization on C_t . To put it in familiar terminology, there is no (coherent) posterior so absurd that no prior could yield it by conditionalization.

²⁴The example, and the point made with it, is from Williamson (2000).

alization we should rationally remain certain that there is one red ball in the urn. Thus, the constraint on rational belief implicit in subjective Bayesianism is also deficient, for although it may well be rational for us to believe that there is one red ball in the urn at the beginning of the experiment, this belief certainly ceases to be rational by the end of it.

Finally, there is the issue that subjective Bayesians are silent on what counts as *evidence*. Above, we assumed various things about it: for instance, that in the experiment the results of each draw count as a piece of evidence, and that your seeing a red and a black ball going into the urn also count as (in this case, temporary) evidence. But these assumptions are exogenous to Bayesianism itself, which is simply silent about what evidence is. Thus, even the implicit constraint on rational belief is empty until we add to subjective Bayesianism a theory of evidence.

7. Knowledge-Based Objective Bayesianism

In giving his theory of “evidential probability,” Williamson seeks to construct a version of Bayesianism that is free of the problems that beset its subjective variety. Williamson’s version of objective Bayesianism, knowledge-based objective Bayesianism, can be presented in the form of three theses.

First, there is a unique evidential probability function, which we will call C_u (for “ur-prior”). This function represents two important things. First, it represents the intrinsic plausibility of any proposition before any evidence is brought to bear on it. For any proposition p , this is given by $C_u(p)$. Second, C_u encodes the bearing that any possible piece of evidence has on any proposition. For any body of evidence E and any proposition p , this is given by $C_u(p|E)$. Notice that this brand of objective Bayesianism is at the opposite end of the spectrum from subjective Bayesianism: subjective Bayesians allow *any* initial credence function, as long as it is probabilistically coherent, whereas Williamson allows only one. There is obviously room for intermediate positions, such as those required by the convergence theorems discussed above.²⁵ Notice too that Williamson explicitly rejects Carnap’s program of figuring out what C_u is on purely syntactic grounds.

Second, Williamson offers the following claim about what evidence a subject has at a time: exactly those propositions which the subject knows at a time. Notice that this means that falsehoods (which cannot be known) cannot ever be part of a subject’s evidence.

Third, and finally, Williamson claims that, if E is the conjunction of all the propositions you know at t , then the credence it is rational for you to assign to any proposition p at t is given by $C_u(p|E)$. This goes beyond merely saying that $C_u(p|E)$ encodes the bearing that E has on p , as we did in the first thesis. It is important to be clear on what the difference is between this claim and the subjective Bayesian requirement of conditionalization. We can get clear on that difference by relying once again on

²⁵For an objection to such moderate positions, see Horowitz (2013).

the experiment with the black and red ball on the urn, which shows that subjective Bayesians cannot accommodate the phenomenon of knowing more by knowing less. According to the subjective Bayesian, your credences at the start of the experiment can be represented by some function C_{t_1} , with the only proviso that it be a probability function (and, if you had any other credence function before it, that it be the result of conditionalizing that previous function on whatever your evidence is at the time). When you see that a red and a black ball are put on the urn (let us call this proposition E), you should change your credences so that, for any proposition p , they are now representable by $C_{t_2}(p) = C_{t_1}(p|E)$.²⁶ This means, in particular, that you should now assign credence 1 to E , which in turn means that you should assign credence 0 to the proposition that both of the balls in the urn are black. Now, once you accumulate as evidence the results of the 10,000 drawings (let's call the proposition which is the conjunction of all those results E'), your credences for any proposition p should now be given by $C_{t_3}(p) = C_{t_2}(p|E')$. Given the cumulative nature of conditionalization, this is also equal to $C_{t_1}(p|E \wedge E')$, which means that you should still assign credence 0 to the proposition that there are two black balls in the urn. On the other hand, according to Williamson's picture, your credences at t_2 (when you see that a black and a red ball are put into the urn) should be given by conditionalizing C_u on whatever you know at that time. Making the simplifying assumption that all you know at t_2 is E , this means that your credences at that time for any proposition p should be given by $C_{t_2}(p) = C_u(p|E)$. For Williamson as well as for the subjective Bayesian, then, at t_2 you should assign credence 1 to the proposition that one of the balls in the urn is not black. But, at t_3 , after you see the results of all 10,000 drawings, arguably you no longer know that one of the balls in the urn is not black. So, your credences at t_3 should *not* be given by conditionalizing on C_{t_2} , but rather by conditionalizing on C_u given what you now know, which we'll call E' : $C_{t_3}(p) = C_u(p|E')$. This means that Williamson's version of how you should update your credences is not cumulative. You fix first what you know at a time, and then you go back to C_u and conditionalize it on what you know.

We can see now how knowledge-based objective Bayesianism answers all the problems for subjective Bayesianism canvassed in the previous section. It answers the subjectivity objection by claiming that there is a unique initial credence function. It answers the objection that subjective Bayesians are silent about evidence by equating the evidence a subject has at a time with what the subject knows at that time. And it answers the objections to conditionalization by replacing that requirement with a non-cumulative account of updating.

Knowledge-based objective Bayesianism contains an explicit commitment about the rationality of credences in its positing of a unique evidential probability function, which determines both which credences it is rational to assign to propositions in the absence of evidence as well as how evidence determines changes in those credences. By themselves, neither knowledge-based objective Bayesianism nor $E = K$ contain any commitments about the rationality of beliefs, except perhaps for the commitment that

²⁶This supplements subjective Bayesianism with substantive assumptions about what your evidence is, but that this supplementation is needed is a deficit of subjective Bayesianism, not an objection to the supplementation.

you rationally believe everything that is part of your evidence. Recently, however, Williamson (forthcoming) himself added another plank to his knowledge-first epistemology. Following Sutton (2005), this recent plank can be seen as amending the slogan to “knowledge first—and last.” The idea is that what a subject knows at a time not only is identical with the subject’s evidence at that time, but also with what the subject is rational in believing at that time.²⁷

8. Knowledge-based Decision Theory

How would a decision theory based on knowledge-based objective Bayesianism model the decision problem with which we began? There are two important stages in that decision problem: one stage represents the problem before your friend tells you that he will go to the movies if it rains and to the park otherwise, and the other stage comes after your friend tells you that. Let us assume that at the first stage you have no idea where your friend will be—which, setting doubts about principle of indifference aside, we represent by your having .5 credence that he will go to the movies and .5 that he will go to the park.²⁸ Let us also assume that the weather report rationalizes a .9 credence that it will rain. Finally, let us assume (as before) that your utility function assigns 9 to meeting your friend, 8 to being at the park without your friend, and 1 to being at the movies without him. If so, the first stage of the problem can be represented in the following matrix:

Table 3: First stage

	$C(\text{Rain} \wedge P) = .45$	$C(\text{Rain} \wedge M) = .45$	$C(\text{No rain} \wedge P) = .05$	$C(\text{No rain} \wedge M) = .05$
Park	9	8	9	8
Movies	1	9	1	9

The expected utility of going to the park is 8.5, and the expected utility of going to the movies is 5. The result is that you are rationally required to go to the park, which makes sense in light of the fact that you have no idea where your friend will be, and you prefer the park to the movies otherwise. The second stage of the problem occurs after your friend tells you that he will be at the movies if it rains and at the park otherwise. In the original case, you acquire knowledge through your friend’s testimony. Thus, according to Williamson’s picture, it is now part of your evidence that your friend will be at the park if it doesn’t rain, and at the movies otherwise. You should therefore conditionalize C_u on this evidence, with the result that you should zero-out your credences in two of the four states that you had considered possible beforehand: the state where it rains and your friend goes to the park, and the state where it doesn’t rain and your friend goes to the movies. Your credences in the remaining two states should retain

²⁷Williamson has argued that this is not an additional plank, but that it follows from $E = K$ —see Williamson (2013a), Williamson (2013b), Williamson (forthcoming), and cf. Cohen and Comesaña (2013a), Cohen and Comesaña (2013b), and Cohen and Comesaña (forthcoming).

²⁸I also assume that you are certain that your friend is at the movies if and only if he is not at the park.

the same proportions they had before, with the result that the matrix describing is the same one that we used before:

Table 4:

	$Cr(\text{Rain} \wedge \text{M}) = .9$	$Cr(\text{No rain} \wedge \text{P}) = .1$
Park	8	9
Movies	9	1

As before, the expected utility of going to the park is 8.1, whereas the expected utility of going to the movies is 8.2. This means that you are rationally required to go to the movies. I leave consideration of whether this is the right result for later.

How would Williamson model the two variations of the case? The first stage of the problem is common to all variations, so let us look at how Williamson would model the second stage of our two variations. The first of those variations is where it probably doesn't rain. That variation doesn't need a new matrix, for its second stage is represented by the same matrix as in the original case. You rationally went to the movies and something unexpected happened, whereby the act with the maximal expected utility turned out to be the one with the least actual utility. For Williamson as well as for anyone else, such is the hard life of the rational agent.

But what about the second variation? There, it does rain (as expected), but your friend ends up not going to the movies for whatever reason. Williamson's analysis of this case is going to be different from his analysis of the original case, for now you do not know that your friend will be at the park if and only if it doesn't rain (if only because that is false). Therefore, you should not assign credence 0 to the state where it rains and your friend goes to the park, or to the state where it doesn't rain and your friend goes to the movies. What credences, then, should you now assign to those states? The case that we are concerned with is one where the belief in question was acquired via your friend's testimony. It matters whether we view testimonially acquired beliefs as inferential or not. If we do, then Williamson has a natural answer to the question about which credence you should assign to the whereabouts of your friend in the case where what he tells you is false—although, as we will see, this natural answer is highly problematic. If we see testimony as giving rise to non-inferentially justified beliefs (as it is natural to think of experience), then it is not quite clear what Williamson can say. I start with the easy case first.

If we are thinking of testimonially acquired beliefs as inferentially justified, then presumably they are inferred from the testimony itself. We can assume that you do know that your friend told you that he would go the movies if it rains, and to the park otherwise. In that case, the credence that you should assign to the states in question is whatever the conditional probability of those states is given your friend's testimony, according to C_u . Presumably, again, that will be a high-ish credence. But exactly how high? The precise number turns out to matter. For, if we set it at, say, .99, then the resulting matrix is the following:

Whereas, if we set the conditional probability in question to "only" .9, the resulting matrix is the following:

Table 5: No knowledge - .99

	$C(\text{Rain} \wedge P) = .009$	$C(\text{Rain} \wedge M) = .891$	$C(\text{No rain} \wedge P) = .099$	$C(\text{No rain} \wedge M) = .001$
Park	9	8	9	8
Movies	1	9	1	9

Table 6: No knowledge .9

	$Cr(\text{Rain} \wedge C) = .09$	$Cr(\text{Rain} \wedge S) = .81$	$Cr(\text{No rain} \wedge C) = .09$	$Cr(\text{No rain} \wedge S) = .01$
Park	9	8	9	8
Movies	1	9	1	9

In the first case, where you take your friend to be super-reliable, the expected utility of going to the park is 8.108 whereas the expected utility of going to the movies is 8.136, with the result that the rational option is to go to the movies, just as when the case involves knowledge on your part. In the second case, however, the expected utility of going to the park is now 8.18, whereas the utility of going to the movies is 7.56—which means that you should now go to the park instead.

But what if we think of testimony as giving rise to non-inferentially justified belief? More generally, regardless of our take on justification by testimony, what can Williamson say about the bad counterparts of non-inferentially acquired knowledge? Take a case where you have non-inferential knowledge that p , and then compare it with a case just like it except that p is false. What, according to Williamson, should your doxastic attitude towards p be then? This case is harder for Williamson to answer than the inferential one, for there is no previous doxastic state to fall back into to determine what your attitude towards p should be. You shouldn't, according to Williamson, believe that p , because you do not know that p . You shouldn't either disbelieve that p , because you don't know that *not-p* either. The only answer left open in the realm of coarse-grained attitudes is that you should suspend judgment regarding p . But there are two problems with this answer. First, this is nothing like paradigmatic cases where suspension of judgment is justified.²⁹ Second, and more importantly in the context of the connections between epistemology and decision theory, which actions would such suspension rationalize? We need to know not only which coarse-grained attitude you should take towards p , but also which credence you should assign it. In the inferential case, the answer was simple: whatever credence is the result of conditionalizing C_u on what you know. But, we are assuming, what you know is evidentially irrelevant to p . So, there are only two options left open to Williamson. First, he may say that you should assign to p whatever credence you should have assigned to it in the absence of your apparent knowledge that p . This is, in essence, to treat the non-evidential case as if it were a null-version of the evidential case. Second, he may say that you now find yourself in a rational dilemma, where no credence assignment to p is rational.

So far, I have just explained how a Williamson-style objective Bayesian would model

²⁹There is some textual evidence that Williamson himself would not go for suspension of judgment—see the admittedly somewhat opaque remarks about Pyrrhonian skepticism in Williamson (forthcoming).

our case and its variations. Next, I explain why those modeling assumptions lead to incorrect results.

9. Is all Knowledge Evidence?

The general question that we are interested in concerns subjects who are mistaken counterparts of subjects who know. It is important to emphasize that this doesn't just happen in skeptical scenarios. It happens also when your friend goes to the park even though it rains, and when your son takes the car keys from the table while you are showering, and in a thousand different cases every day. A preoccupation with such cases cannot be dismissed as a scholastic concern in bizarre scenarios—not that there's anything wrong with that. The question is which credence should these subjects assign to the propositions known by their counterparts? As we just saw, there are two interestingly different cases: the evidential and the non-evidential one. It's clear how a proponent of knowledge-based objective Bayesianism should treat the evidential case—but it leads to absurd consequences. It's not at all clear how such a proponent should treat the non-evidential case—but, again, none of the possible options are attractive. In this section I start with a problem arising from the evidential case.

Let's go back to our case and consider what happens when we treat testimony as issuing in inferentially justified beliefs. In the previous section we saw that a proponent of knowledge-based decision theory can say that the credence you should assign an inferentially justified proposition when you do not know it is whatever conditional probability it has on the evidence on which it is based. In the previous section we saw two interestingly different cases depending on what exactly this conditional probability is. I'll draw some consequences from this difference momentarily, but for now I want to highlight another problem for knowledge-based decision theory.

In the previous section I was implicitly assuming that it is possible to have inferentially acquired knowledge even when the conditional probability of the proposition known given the evidence on which it is based is less than 1. Of course, given $E = K$, everything a subject knows at a time has probability 1 on his total evidence at that time. But we must not confuse this (trivial, given $E = K$) fact with the substantial, and false, claim that inferential knowledge of a proposition p is possible only when the rational credence in p is raised to 1 by the evidence on which it is based. Let E be the proposition that your friend said that he will be at the movies if it rains and at the park otherwise (and call this proposition H). If on the basis of his testimony you know at $t1$ that what he said is true, then your evidence at $t1$ includes H , and so $Cr_{t1}(H|E) = 1$, but this doesn't mean that $Cr_u(H|E) = 1$.³⁰ Call this kind of knowledge, inferential knowledge based on non-entailing evidence, inductive knowledge.

³⁰I have argued that when you are inferentially justified in believing a proposition you are justified in believing propositions which entail it—see Comesaña (2013), Comesaña (forthcoming) and Comesaña (2015). But I wasn't assuming there that evidence is what you should conditionalize on. If we do make that assumption (as we are doing here), the conclusion of those papers can be re-stated as the claim that whenever a subject is inferentially justified in assigning a credence n to a proposition, he is non-inferentially justified in assigning credence at least n to propositions which entail it.

So, $E = K$ leaves it open that inductive knowledge is possible. It is not clear whether Williamson himself thinks that inductive knowledge is possible.³¹ But I think it is fair to say that most epistemologists will take it as a reductio of a view that it doesn't allow for the possibility of inductive knowledge in this sense. Let us therefore assume that the proponent of $E = K$ allows for inductive knowledge.

The combination of $E = K$ with inductive knowledge, however, leads to disaster. For assume that $Cr_u(H|E) < 1$, and that a subject S knows that E at t and on that basis comes to know that H . In that case, $Cr_t(H) = 1$. But this is the wrong credence for the subject to have in H —this is irrational overconfidence. Of course, as we said before, if $E = K$ then the evidence that S has at t rationalizes a credence of 1 in H , but (modulo $E = K$) the rational credences in an inferred proposition at a time are not the credences rationalized by the evidence had at that time, but rather by the evidence on the basis of which the proposition is inferred. Moreover, this irrational credence will inevitably lead to further irrational credences, for there will certainly be propositions whose credence is changed by conditionalizing on inductive knowledge. Thus, there will inevitably be some proposition H' such that $Cr_u(H'|E) \neq Cr_u(H'|H)$, and $E = K$ has the wrong consequence that the rational credence in H' , after one comes to inferentially know H , are given by the latter.

These irrational credences can have downstream effects on action. Suppose that $Cr_u(H|E) = .9$ —where, remember, E is your friend's testimony and H its content. In that case, it is rational for you to go to the park—you think your friend is reliable, but not enough to overcome your preference for the park. However, suppose that you know that what your friend said is true. In that case, according to $E = K$ you should conditionalize on E . But, if you conditionalize on E , going to the movies maximizes expected utility. Therefore, if we combine $E = K$ with the possibility of inductive knowledge, not only do we get absurd epistemological consequences, we also get absurd consequences regarding rational action.³²

Williamson himself has come tantalizingly close to this problem. Attacking an alternative to $E = K$, he says the following:

If evidence required only justified true belief, or some other good cognitive status short of knowledge, then a critical mass of evidence could set off a kind of chain reaction. Our known evidence justifies belief in various true hypotheses; they would count as evidence too, so this larger evidence set would justify belief in still more true hypotheses, which would in turn count as further evidence. ...The result would be very different from our present conception of evidence. (Williamson (2000), p. 201)

³¹The material surrounding the passage from *Knowledge and Its Limits* quoted later in this section is interestingly ambiguous on this issue.

³²When conditionalizing on inductive knowledge rationalizes irrational actions, the consequences can be avoided by holding on to something like the knowledge-to-action principle defended by Fantl and McGrath—see Fantl and McGrath (2002), Fantl and McGrath (2007) and Fantl and McGrath (2009). Williamson himself thinks that the knowledge-to-action principle is false, proposing instead that the cases in question can be taken care of by noticing a failure of knowledge to iterate—see Williamson (2005). In any case, even when inductive knowledge does not give rise to irrational actions, conditionalizing on it leads to irrational credences.

As Alvin Goldman has seen, the most puzzling feature of this passage is that the criticism presented in it applies pretty obviously to $E = K$ as well (assuming the possibility of inductive knowledge):

(...) I am puzzled that Williamson presents this argument against the JTB [justified true belief] account of evidence, because the contemplated chain reaction seems *almost* as threatening under the $E = K$ construal. Clearly, $E = K$ does not restrict evidence to non-inferential knowledge. So if one's 'basic' evidence justifies belief in various true hypotheses that are also known, a very similar chain reaction looms. (Goldman (2009), p. 88.)³³

10. The Irrelevance of Truth

The problem for $E = K$ presented in the previous section can be bypassed by moving to $E = \text{basic } K$. That is to say, instead of identifying evidence with all of the subject's knowledge, one could identify it with that subset of his knowledge which was not inferentially acquired. In this section, I present more serious problems for any factive account of evidence.

Consider again the proposition that your friend will go to the park if it doesn't rain, and to the movies otherwise. According to Williamson, whereas in the original version of the case you are rational in believing this proposition (and, not coincidentally, also rational in assigning it credence 1), in the second variation of the case you aren't.³⁴ One way to put the objection that I want to develop in this section is to say that the only relevant difference between the original version and the second variation of the case is in the mere truth of the proposition in question, and this is an epistemically insignificant difference. Now, Williamson will not accept this characterization of the difference, for in the original version you know the proposition. Given that Williamson takes knowledge to be a bona fide mental state, he can argue that the mere truth of the proposition makes an epistemically relevant difference.

It can still be argued that the only relevant difference between the cases is in the mere truth of the proposition in question, even giving Williamson his controversial thesis about knowledge being a mental state. For knowledge entails truth. If so, then (arguably) there are cases where the only difference between a case of knowledge and a case of lack of knowledge (the only difference in whether this mental state is instantiated, if Williamson is right) is the truth of the proposition in question.³⁵ So, if the mere truth of the proposition in question cannot make a difference as to what credence it is

³³In his replies on the same volume, Williamson does not address this argument.

³⁴Another objection to knowledge-first decision theory can be made based on Gettier cases—see Brown (2008).

³⁵See Comesaña (2005). I would now take back some of what I said in that paper, namely, the part about warrant (understood as whatever turns true belief into knowledge) being epistemologically relevant because truth is. I was moved to say this in the dialectical context, but I now think that it is pretty obvious that truth (by itself) is not epistemologically relevant, and probably neither is warrant.

rational to assign to it, then at least sometimes whether you know a proposition cannot make a difference as to what credence it is rational to assign to it.

One philosopher's Modus Ponens, however, is another's Modus Tollens. So Williamson could argue that, given that knowledge can make a difference as to what beliefs and credences it may be rational to have, and given that sometimes the only difference between a case of knowledge and one of ignorance is the truth of the proposition in question, sometimes the mere truth of a proposition can make a difference as to what beliefs and credences it is rational to have. But the Williamsonian Modus Ponens faces serious problems that do not affect my Modus Tollens.

Consider the following very plausible assumption: that any condition that can make a difference in what beliefs and credences it is rational to have must be a condition whose obtaining is rationally believable in the circumstances in which it makes a difference. This is a very plausible and minimalist cashing out of the metaphor that rational norms should be "action guiding." Notice that the assumption does not say that the obtaining of the condition must be actually rationally believed, only that it should be possible for it to be rationally believed in the circumstances. On the other hand, the assumption is not the extremely weak claim that the obtaining of the condition should be rationally believable in some circumstance or other—this would make it a trivial condition, for almost all conditions are rationally believable in some circumstance or other (barring unknowable conditions, of course).

In the variation of the case that we are considering, you of course don't know that what your friend told you is true, because it is false. But you also don't know that you don't know it. And this is not just because you don't believe that you don't know it, but because your epistemic position with respect to the proposition that you don't know it is not strong enough. Your evidence makes it highly likely that you do know it, and if your evidence makes it highly likely that p , then it is irrational for you to believe that $not-p$, and hence you can't know that $not-p$. Therefore, it is irrational for you to believe that you don't know that what your friend says is true, and so that fact cannot make a difference as to what credences it is rational for you to assign.

The requirement of the rational believability of the conditions which determine rational credences does not amount to a requirement of infallibility. Williamson has rightly pointed out that just because we can make mistakes in following a rule that doesn't mean that we are following a different rule.³⁶ So, the mere fact that we can be mistaken about what we know or don't know doesn't show that we cannot follow the rule that we should adjust our credences to what we know. The objection is not that we cannot follow the rule, but that there are circumstances where it would be irrational to follow it. And if it would be irrational to follow the rule, then it cannot be rational to form the attitudes the rule requires—at the very least, it cannot be rational to form those attitudes *by* following the rule.

On a knowledge-based decision theory, not only the rationality of credences, but also the rationality of actions is determined by sometimes unknowable (and not rationally believable) conditions. The conditional probability of the truth of what your friend

³⁶Williamson (2000), pp. 191-2.

says given that he says it determines, in the case where you do not know what your friend tells you, whether you should go to the park (if, say, the probability in question is .9) or to the movies (if, say, the probability in question is .99). Keeping preferences fixed, and given that there is a unique evidential probability function, which actions it is rational for you to perform at a time supervenes on which propositions you know at that time. But, as I just argued, sometimes the difference between knowing and not knowing a proposition will display the kind of epistemic asymmetry familiar from consideration of skeptical scenarios: if you know, you can know that it is not the case that you do not know, but if you do not know, then you cannot know that it is not the case that you know, because it is irrational for you to believe that you do not know. That is what happens with whether your friend is telling the truth. And just as such conditions cannot affect which credences it is rational to assign, they cannot affect either which actions are rational to perform.

That last example depended on an inferentialist account of testimonial justification. But things are not better in the non-inferential case. Take a subject who non-evidentially knows that p , and consider a counterpart whose situation differs from his only in that p is false (and, of course, whatever else supervenes on it). Which credence should that counterpart assign to p ? In the previous section, I argued that there are two potential answers to this question. First, it may be held that he should assign whatever credence is the result of conditionalizing C_u on what the subject knows. This answer in effect treats the non-inferential case as inferential, and it runs into the same problems that we just canvassed. Second, it could be argued that there is no credence that is rational for this subject to assign to p , and so, for those cases where which action is rational depends on which credence p is assigned, there is no action which it is rational for the subject to perform.

My own view that *ex-post* rational action requires *ex-post* rational beliefs and credences countenances the existence of rational dilemmas. As explained in section 4, when the subject does not have *ex-post* rational beliefs and credences there is no action which is *ex-post* rational for the subject to perform. But I argued in that section that the kinds of dilemmas admitted by my view have three characteristics which makes them acceptable: they are only *ex-post*, and not also *ex-ante*, dilemmas; they are self-imposed; and they are recoverable. By contrast, the kind of dilemmas that a knowledge-first view would be committed to have none of these properties, and so they are deeply problematic.

First, the kind of dilemma posited by the knowledge-first view is *ex-ante* as well as *ex-post*. The problem is not that the subject assigns a credence (whether the rational one or not) for the wrong reasons: the problem is that there is no rational credence to assign.³⁷ And so it is not just that whatever action the subject performs will not be based on an *ex-post* rational credence assignment, but, in addition, there is no *ex-ante* rational action for the subject to perform. Second, the kind of dilemma in question is not self-imposed. The position the subject is in is due exclusively to the fact that, were

³⁷The view that there is a rational credence for the subject to assign is the view that treats the non-inferential case as if it were inferential, and as we said it runs into the same problems as the knowledge-first treatment of the evidential case.

the proposition in question true, he would know it. Third, and relatedly, the dilemma is non-recoverable. There is nothing the subject could do, and no doxastic attitude he could adopt, that will take him out of the dilemmatic situation.

Although the existence of merely *ex-post*, self-imposed and recoverable dilemmas is not implausible, it is implausible that there are *ex-ante*, world-imposed and non-recoverable rational dilemmas. Notice, moreover, that in the knowledge-first view, these dilemmas will be absurdly widespread: they exist every time a subject believes a false proposition that he would know if it were true.

In this section I have argued against knowledge-based decision theory. My argument focused on cases where a subject believes a false proposition that he would know if it were true. We can usefully distinguish two such cases: evidential and non-evidential ones. Evidential cases of that kind are a sub-class of cases of misleading evidence. Knowledge-based decision theory does accept the existence of cases of misleading evidence—they happen whenever the subject has some evidence E and, for some false proposition p , $C_u(p|E) > C_u(p)$. If, moreover, $C_u(p|E)$ is high enough for the subject to know p , if it were true, then we have an evidential case of the sort I am discussing here. The knowledge-based decision theorist is forced to say that, in those cases, the subject should assign credence $C_u(p|E)$ to p , even though if p were true he should assign it credence 1. In those cases, the subject is not in a position to know either that p is false or that he doesn't know that p , and so which credence he should assign depends on an unknowable condition. Moreover, because which action is rational for the subject to perform may depend on which credence the subject assigns to the proposition in question, which action is rational also depends for such subjects on an unknowable condition. The non-evidential case can be treated (by the knowledge-based decision theorists) just as the evidential case. If so, the same problems arise. Alternatively, it can be treated as giving rise to rational dilemmas. The dilemmas in question would be *ex-ante*, world-imposed and non-recoverable, and so their existence would be deeply problematic.

My own view is that, for these reasons, we should prefer rationality-based decision theory to knowledge-based decision theory. Before presenting that theory, there is a further issue that I need to address: in the next section, I examine the “excuses maneuver,” a defense of knowledge-based decision theory according to which false beliefs can be excusable, but not rational.

11. Excuses, excuses

Recently, Williamson has answered the charge that his view has the wrong consequences for the rationality of credences by appealing to a distinction between rationality and excuses, and the same move has been previously sketched by Hawthorne and Stanley.³⁸ Williamson himself doesn't apply the distinction to the rationality of action, but if his defense worked for the case of credences it would equally work for the

³⁸See Williamson (forthcoming) and Hawthorne and Stanley (2008). As I said in the Introduction, similar defensive maneuvers can be implemented by distinguishing between epistemic justification and rationality.

case of action. In this section I go along with Williamson in thinking that it is rational to assign credence 1 to a proposition you know to be true.

Williamson's idea is that although it is not rational to assign credence 1 to the proposition that your friend is right when you do not know it, you are excused for doing it. This is so because a rational person in your position would assign credence 1 to that proposition, and if you do what a rational person would do in your position, then you are at least excused for doing so. The crucial question for Williamson, then, is: why are you not rational in assigning credence 1 to that proposition, given that it is the credence that a rational person would assign in your position? In answer to this question, Williamson argues that, in general, it is false that the rational credence to assign to a proposition in a given situation just is the credence that a rational person would assign it in that situation. Williamson's argument appeals to a case. Suppose that there is a brain-scrambling mechanism that interferes with the subject's short-term memory. Under the effects of such a brain-scrambler, a rational person might believe anything—including propositions it would not be rational to believe.

It is not clear that Williamson's example shows what he takes it to show. It is crucial for his example to do the work Williamson needs it to do that the subject in question retains his rationality under the influence of the brain-scrambler, but it is not at all obvious that he does so. But I will not press that point here, because there are more serious problems with Williamson's position.

There are two objections to Williamson's treatment of the case where you do not know that your friend is telling the truth. The first objection is that it has the consequence that you would not be rational to assign credence 1 to that proposition. Against this objection, the distinction between rationality and excuses (and the underlying distinction between the rational credence and the credence a rational subject would assign) makes sense: although it would indeed be irrational to assign credence 1 to that proposition, it would be excusable to do so, and perhaps the objector has simply conflated rationality with excusability. This first objection accuses Williamson of omitting to say that you are rational in assigning credence 1 to the proposition that your friend is telling the truth when you do not know that proposition. The second objection is different: it accuses Williamson of incorrectly saying that you are rational to assign a credence less than 1 to the proposition in question. The distinction between rationality and excusability has no obvious application against this objection.

Precisely because the rational subject would assign to a proposition he takes himself to know the same credence he would assign it if he did know it, the ideal Williamsonian subject is literally unintelligible. The ideal Williamsonian subject assigns to every proposition the credence dictated by C_u conditionalized on everything the subject knows at a time. Thus, this ideal subject assigns credence less than 1 to the proposition that his friend is telling the truth in the second variation of our case. Imagine a dialogue with him:³⁹

³⁹Perhaps you are a skeptic, or disagree with Williamson about the role of knowledge in the update on credences. If so, imagine a dialogue with an ideal Williamsonian subject who is going to the movies because he knows his friend is telling the truth.

You: Why are you going to the park? It's probably going to rain, and your friend told you that he would go to the movies if it rains.

Ideal Williamsonian Subject: Yes, but still, the evidentially expected utility of going to the park is higher than that of going to the movies, given what I know.

You: But it isn't if you know that your friend is telling the truth.

IWS: Oh, but I don't know that.

You: Really? Is your friend a habitual liar?

IWS: Oh no, he is a paragon of honesty.

You: So, is there any other reason to think that what he says is false?

IWS: Not really.

You: So why on Earth do you not take him at his word and go to the movies?

IWS: . . .

There are situations where we cannot know that we do not know. In those situations, it is rational to act on what is rational to think we know, not on what we know. Williamson's model cannot deliver that result.

12. Rationality-Based Decision Theory

For all its knowledge-first rhetoric, at the core of Williamson's epistemology sits a profoundly rationality-first thesis: the existence and uniqueness of C_u , the evidential probability function. For remember that we can see this function as encoding two important things: first, the credence that it is rational to assign to any proposition in the absence of any evidence for it (encoded, for any proposition p , in $C_u(p)$); second, the bearing of any body of evidence E on any proposition p (encoded, for any E and p , in $C_u(p|E)$). Williamson does superimpose to this rationality basis a knowledge-first account of evidence, but this is an independent and optional component of the total view. In this section I develop a view which takes on the rationality basis but leaves behind the optional knowledge-first component.⁴⁰

⁴⁰The theory of evidence defended in this section has been proposed by Goldman (2009) as superior to $E = K$, but not on the basis of the decision-theoretic arguments here presented. Weatherson (ms) argues against $E = K$ and for the identification of evidence with the outputs of reliable Fodorian modules, allowing explicitly for false evidence. Surprisingly, Weatherson (2012) defends knowledge-based decision theory and considers justified-true-belief-based decision theory, but not a plain rationality-based (or justification-based) theory. He thinks that the property of being a justified true belief is less natural than the property of being a known proposition, and then seems to conclude that any property weaker than knowledge will be less natural. That is a plainly invalid argument, however, and so I hesitate to ascribe it to Weatherson. Dutant (forthcoming) examines and rejects knowledge-based decision theory, and some of his reasons are related to the ones I develop in this paper. However, he goes on to defend a supported-by-knowledge decision theory. Dutant's theory and my own will differ in cases of basic rationality—that is to say, cases where a belief is rational but not based on any other beliefs. My theory, but not Dutant's, allows such beliefs to have a say in the determination of the decision matrix. I see this as an advantage of my view over Dutant's.

My view can be presented in the form of three theses, the first two of which are shared with Williamson's view.

Those two first theses are that there is a unique evidential probability function C_u and that the credence that a subject should rationally assign to any proposition p are those determined by $C_u(p|E)$, where E is the conjunction of all the evidence the subject has at that time.

The third thesis is that the evidence that a subject has at a time is composed by all and only those propositions that the subject is basically rational in believing at that time. I am here replacing Williamson's $E = K$ with $E = \text{basic R}$. As argued above, there are reasons for even the knowledge-firster to move from $E = K$ to $E = \text{basic K}$, and those reasons favor the identification of evidence with basically (non-evidentially) rational belief over its identification with rational belief simpliciter. The view leaves it open that there are evidentially justified beliefs, but they are not part of the subject's evidence.

Williamson's objective Bayesianism has several advantages over subjective Bayesianism, but also suffers from serious problems of its own. The advantages are that it avoids the subjectivity of subjective Bayesianism, it avoids the problems with conditionalization as an updating rule, and it provides a theory of evidence. The problems are that it allows for the rationality of both credences and actions to depend on conditions whose obtaining cannot be rationally believed. It also has the wrong result regarding some cases, but rather than additional problems these extensional mistakes are a symptom of the deeper issues just mentioned.

My position inherits all the advantages of Williamson's, while avoiding the disadvantages. My first two theses, shared with Williamson, avoid the subjectivity of subjective Bayesianism and its problems with conditionalization. My third thesis replaces Williamson's $E = K$ with $E = \text{basic R}$, which avoids the consequence that which credences it is rational to assign might depend on unknowable conditions. This is so because although the mere truth of the proposition in question might be the only relevant difference between a case of knowledge and one of ignorance, this does not happen with the difference between a rational and an irrational belief, given that false beliefs (pace Williamson) can well be rational.

My view has all the right consequences for the cases discussed in this paper. What exactly those consequences are depends on whether we take testimony to issue in basic justified beliefs or instead in inferentially justified beliefs. If we take testimony as issuing in inferentially justified beliefs, then it is never part of your evidence, in my view, that what your friend said is true. What is part of your evidence is that your friend said that it is true.⁴¹ This evidence might well justify you in assigning a high degree of credence to the proposition which is the content of your friend's testimony—indeed, given the right background, it might even justify you in assigning credence 1 to that content. In any case, the credence you are justified in assigning to that proposition depends on what the correct evidential relation is and what else is part of your evidence.

⁴¹I'm taking this a non-inferentially justified belief. If you disagree, substitute your favorite candidate. My view does assume that there are some non-inferentially justified beliefs, and to this extent it is a foundationalist view.

Whichever credence is thus justified, you are rational in doing whatever maximizes expected utility relative to that credence.

Suppose instead that we take testimony as issuing in basically justified belief (or, if you think this is the wrong view about testimony, just take the case as a proxy for your favorite source of basic justification). In that case, whether your friend is telling the truth or not, you are justified in believing him given the circumstances, and so it is part of your evidence that what he said is true. Therefore, it is rational for you to go to the movies regardless of whether what he said is true.

Williamson himself has briefly presented two arguments against non-factive conceptions of evidence such as mine, but I don't have much to say about them.⁴² The first argument is that if evidence can be false then evidence can rule out true propositions. To that, I can only say that that is indeed the view. The second one is that the view that evidence can be false gains most of its appeal from the assumption that we have infallible access to our evidence. The reader can verify that such an assumption played no role in my arguments.

13. Conclusion

Recall that my starting point were judgments about the rationality of actions in a variety of cases. From that starting point I drew out epistemological consequences. In particular, I argued against a Williamsonian version of objective Bayesianism and in favor of my own, rationality based alternative to it. But this dialectical order belongs to the context of discovery, to appropriate Reichenbach's useful terminology. In the context of justification, epistemology comes first. Much as Williamson does not have a lot to say about when a proposition is known, I do not have a lot to say (here) about when you are rational to believe a proposition. After we fix what you are basically rational in believing, we then fix which credences it is rational for you to assign to different propositions by conditionalizing on C_u . Finally, that determines, via expected utility maximization, which actions are rational for you to perform.

References

- Brown, Campbell. 2011. "Consequentialize This." *Ethics* 121 (4): 749–71.
- Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Noûs* 42 (2): 167–89.
- Christensen, David. 2004. *Putting Logic in Its Place*. Oxford University Press.
- Cohen, Stewart. forthcoming. "Theorizing About the Epistemic." *Inquiry*.
- Cohen, Stewart, and Juan Comesaña. 2013a. "Williamson on Gettier Cases and Epistemic Logic." *Inquiry* 56 (1): 15–29.

⁴²Williamson (2000), pp. 201–2.

- . 2013b. “Williamson on Gettier Cases in Epistemic Logic and the Knowledge Norm for Rational Belief: A Reply to a Reply to a Reply.” *Inquiry* 56 (4): 400–415.
- . forthcoming. “Being Rational and Being Right.” In *The New Evil Demon*. Oxford University Press.
- Comesaña, Juan. 2005. “Justified Vs. Warranted Perceptual Belief: Resisting Disjunctivism.” *Philosophy and Phenomenological Research* 71 (2): 367–83.
- . 2013. “There Is No Immediate Justification.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup Ernesto Sosa and John Turri. Blackwell.
- . 2015. “Falsehood and Entailment.” *Philosophical Perspectives* 29 (1): 82–94.
- . 2016. “Normative Requirements and Contrary-to-Duty Obligations.” *Journal of Philosophy* 112 (11): 600–626.
- . forthcoming. “On Sharon and Spectre’s Argument Against Closure.” *Philosophical Studies*.
- Dutant, Julien. forthcoming. “Knowledge-Based Decision Theory and the New Evil Demon.” In *The New Evil Demon*, edited by Julien Dutant and Fabian Dorsch. Oxford University Press.
- Fantl, Jeremy, and Matthew McGrath. 2002. “Evidence, Pragmatics, and Justification.” *Philosophical Review* 112 (1): 47–67.
- . 2007. “On Pragmatic Encroachment in Epistemology.” *Philosophy and Phenomenological Research* 75 (3): 558–89.
- . 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- Gibbard, Allan, and William Harper. 1978. “Counterfactuals and Two Kinds of Expected Utility.” In *Foundations and Applications of Decision Theory*, edited by James L. Leach Clifford Alan Hooker and Edward Francis McClennan, 125–62. Reidel.
- Goldman, Alvin. 1979. “What Is Justified Belief?” In *Justification and Knowledge*, edited by George Pappas, 1–23. Dordrecht: Reidel.
- . 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- . 2009. “Williamson on Knowledge and Evidence.” In *Williamson on Knowledge*, edited by Patrick Greenough and Duncan Pritchard, 72–91. Oxford University Press.
- Hawthorne, John, and Jason Stanley. 2008. “Knowledge and Action.” *Journal of Philosophy* 105 (10): 571–90.
- Hájek, Alan. 2003. “Conditional Probability Is the Very Guide of Life.” In *Probability Is the Very Guide of Life: The Philosophical Uses of Chance*, edited by Henry Kyburg and Mariam Thalos. Open Court.
- Horowitz, Sophie. 2013. “Immoderately Rational.” *Philosophical Studies* 167 (1): 41–56.
- Jeffrey, Richard. 1965. *The Logic of Decision*. 2nd. McGraw Hill.

- Lasonen-Aarnio, Maria. forthcoming. "Virtuous Failure and Victims of Deceit." In *The New Evil Demon*, edited by Fabian Dorsch and Julien Dutant. Oxford University Press.
- Lehrer, Keith. 1971. "How Reasons Give Us Knowledge, or the Case of the Gypsy Lawyer." *The Journal of Philosophy* 68 (10): 311–13.
- Littlejohn, Clayton. forthcoming. "A Plea for Epistemic Excuses." In *The New Evil Demon*, edited by F. Dorsch and Julien Dutant. Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*. Oxford University Press.
- Savage, Leonard. 1954. *The Foundations of Statistics*. John Wiley; Sons.
- Schellenberg, Susanna. 2013. "Experience and Evidence." *Mind* 122 (487): 699–747.
- Stalnaker, Robert. 1991. "The Problem of Logical Omniscience, I." *Synthese* 89 (3): 425–40.
- . 1999. "The Problem of Logical Omniscience, II." In *Context and Content*, 255–73. Oxford University Press.
- Sutton, Jonathan. 2005. "Stick to What You Know." *Nous* 39 (3): 359–96.
- . 2007. *Without Justification*. Massachusetts Institute of Technology.
- von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- Weatherson, Brian. 2012. "Knowledge, Bets, and Interests." In *Knowledge Ascriptions*, edited by Jessica Brown and Mikkel Gerken, 75–103. Oxford University Press.
- . ms. "E ≠ K."
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.
- . 2005. "Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge." *The Philosophical Quarterly* 55 (219): 213–35.
- . 2013a. "Gettier Cases in Epistemic Logic." *Inquiry* 56 (1): 1–14.
- . 2013b. "Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic." *Inquiry* 56 (1): 77–96.
- . forthcoming. "Model-Building in Philosophy." In *Philosophy's Future: The Problem of Philosophical Progress*, edited by Russell Blackford and Damien Broderick. Blackwell.
- . forthcoming. "Justification, Excuses and Sceptical Scenarios." In *The New Evil Demon*, edited by Julien Dutant and Fabian Dorsch. Oxford University Press.