

AN EXTERNALIST DECISION THEORY FOR A PRAGMATIC EPISTEMOLOGY

Abstract In recent years, some epistemologists have argued that practical factors can make the difference between knowledge and mere true belief. While proponents of this pragmatic thesis have proposed necessary and sufficient conditions for knowledge, it is striking that they have failed to address Gettier cases. As a result, the proposed analyses of knowledge are either lacking or susceptible to counterexamples. Gettier cases are worth reflecting on because they raise foundational questions for the pragmatist. Underlying these challenges is the fact that pragmatic epistemologies not only rely upon normative theories of rational choice but also require externalist standards to rule out epistemic luck. Unfortunately, we lack adequate externalist theories of rational choice. Thus, the main aim of the paper is to address these foundational challenges by offering the outlines of an externalist decision theory. While this task is an ambitious one that I cannot hope to complete, I will offer the outlines of a decision-theoretic framework on which a richer pragmatic epistemology can be developed. My hope is that this framework opens up new avenues of pragmatic exploration.

Central to epistemological orthodoxy is the tenet that only truth-relevant factors, such as the reliability of belief-forming processes or the counterfactual sensitivity of belief, distinguish knowledge from mere true belief. In recent years, however, this orthodoxy has come under attack. A number of epistemologists have argued in favor of **Pragmatic Encroachment**, which is the view that even if we fix all the truth-relevant factors, varying pragmatic factors can make a difference in determining whether or not a subject's true belief counts as knowledge.¹ Let us call any account that embraces pragmatic encroachment, a **Pragmatic Account of Knowledge**.²

The most powerful arguments for pragmatic encroachment appeal to norms that connect knowledge and the practical. Various proponents have converged on some variant of the following principle: S knows that *P* only if S is justified in taking *P* for granted in deliberation.³

This type of principle along with a suitable pair of cases offers an intuitive argument for pragmatic encroachment. For example, consider Catherine who possesses very strong but inconclusive evidence for

¹I am following (Fantl and McGrath 2009) in the narrow use of this term. However, as they note, the term is sometimes used to refer to any account that allows for pragmatic conditions on knowledge. See (Kim 2017) for a discussion of the varieties of pragmatic encroachment.

²Pragmatic accounts like those given by (Stanley 2005), (Hawthorne 2006) (Hawthorne and Stanley 2008), (Fantl and McGrath 2009), and (Weatherson 2012) have been described under the labels 'subject-sensitive invariantism', 'interest-relative invariantism', 'anti-purism', and 'anti-intellectualism'.

³(Fantl and McGrath 2002) comes closest to this formulation. For ease of explanation, 'believing' will be taken to be synonymous with 'taking for granted'.

her true belief that she was born in New York. We may suppose that she has been told so by her honest and reliable parents. Catherine's true belief appears to be a paradigmatic case of knowledge and if she were filling out an employment form, she would be justified in taking her birthplace for granted. But is Catherine justified in taking for granted that she was born in NY regardless of the practical situation she is in? What if Catherine were caught in a sinister scheme where death would be the punishment for falsely answering a question about her birthplace? In this high stakes situation, it would be irrational for Catherine to take her birthplace for granted. After all, given the opportunity, she should go and double-check, gathering as much additional evidence as she can. So the epistemic principle connecting knowledge and deliberation along with our intuitions in these cases entails that even though all the truth-relevant factors remain fixed, Catherine's true belief counts as knowledge in the low stakes scenario but does not count as knowledge in the extreme high stakes scenario. Most pragmatists also claim that the practical situation of the subject is relevant in a very specific way, by determining the epistemic standards that must be met in order to know. So the reason why Catherine fails to know in the high-stakes scenario is because knowing in this scenario requires that an extremely high standard be met.⁴

My aim in this paper is not to evaluate the arguments for and against pragmatic encroachment but rather to consider the pragmatic view in more detail. In particular, I want to explore how a pragmatic account of knowledge, which arises from thinking about the relationship between knowledge and choice, accounts for the incompatibility of knowledge and epistemic luck. To do so, I will first consider how contemporary pragmatic accounts might address Gettier-style counterexamples. As it stands, proponents of pragmatic encroachment have not discussed these cases in any detail.⁵ In section 1, I show that Gettier cases bring to light some problems for and challenges to pragmatic accounts of knowledge. Most importantly, they show that adequate pragmatic theories must appeal to externalist accounts of rational choice. In section 2, I offer one approach to developing an externalist theory of rational choice and on that basis, develop one type of pragmatic account. One upshot of this account is its unique insight into the nature of Gettier cases. In section 4, I conclude by first addressing some potential objections. But more importantly, I show how the proposed framework can be used to explore pragmatic epistemologies

⁴In general, the standards required to be in a position to know do not always vary when the stakes are raised. To come up with general principles, the notion of risk is more appropriate than the notion of practical stakes.

⁵(Douven 2005) is an exception but he simply claims that Gettier cases are under-described cases that fail to identify the relevant practical factors.

by considering some alternatives to the proposed account. The paper takes on and brings together many different topics (e.g. Gettier cases, skepticism, decision theory, and objective measures of evidential strength). Instead of discussing each issue in careful detail, the paper aims to present, in broad brush strokes, the outlines of a new approach to rational choice. My hope is to show that this approach is particularly suitable as a foundation for further explorations of pragmatic epistemologies.

1 The Gettier Challenge for Pragmatists

In recent years, a number of pragmatic accounts of knowledge have been offered. Hawthorne and Stanley propose that “where one’s choice is p -dependent, it is appropriate to treat the proposition that P as a reason for acting if and only if you know that P .”⁶ Fantl and McGrath propose that “you know that P iff: P can be a reason in all three senses – justifying, favoring, and motivating – and it is not a matter of Gettier-like luck that P can be a favoring reason given that it can be a justifying and motivating reason.”⁷ Weatherson proposes that “it is legitimate to write something on the decision table. . . iff the decision maker knows it to be true.”⁸

Unfortunately, each of these proposals is unsatisfactory. At first glance, both Weatherson’s proposal as well as Hawthorne and Stanley’s appear to fall prey to Gettier-style counterexamples.⁹ Consider a simple Gettier case in which Catherine is told by her very honest and reliable parents that I : she will receive a small inheritance when she purchases her first house. Unbeknownst to both her and her parents, her inheritance has been stolen and squandered away by her younger sister. Luckily, a distant relative has set aside an inheritance that will be given to her when she purchases her first home. Intuitively, it seems appropriate for Catherine to treat I as a reason for buying her first house. It also appears legitimate for Catherine to write that she will receive an inheritance as an outcome of buying her first home on the decision table. After all, it’s true and she possesses good evidence in favor of its truth. However, since her true belief has been Gettiered, she fails to know in both instances. So each account appears to fall prey to Gettier-style counterexamples.

⁶p.578, (Hawthorne and Stanley 2008)

⁷p.175, (Fantl and McGrath 2009)

⁸p.77, (Weatherson 2012) Weatherson also proposes that “it is legitimate to leave a possible state of the world off a decision table iff the decision maker knows it not to obtain.”

⁹(Brown 2008) uses Gettier situations as counterexamples to Hawthorne and Stanley’s analysis.

Of course, a similar reply is available for both accounts. Hawthorne and Stanley can respond by simply insisting that Catherine may not treat *I* as a reason for acting, and she may not do so because she doesn't know. Similarly, Weatherston can respond by insisting that Catherine's entry in the decision table is illegitimate because she doesn't know. While these replies preserve the extensional adequacy of the two accounts, they undermine their explanatory power. After all, we have no independent account of when one may treat *I* as a reason or when one may write something on the decision table that explains why we fail to know in Gettier cases. Instead, the lack of knowledge in Gettier cases is used to explain why the proposed conditions for knowing are not met. So we are left with accounts that offer little insight into the nature of knowledge. Instead of gaining understanding about knowledge by appealing to the notion of reasons or a theory of rational decision-making, these accounts provide insight into when one may treat something as a reason or write something on a decision table by considering what we do and do not know.¹⁰

This latter criticism also applies to Fantl and McGrath's proposal. Since the account appeals to an explicit anti-Gettier condition without any explanation of what it means to meet this condition, we are left with the uninformative claim that Gettiered subjects fail to know because they are in a Gettier case. While their account does connect the concepts of knowledge and reason, we are nevertheless left without a clear account of the type of reason knowledge is.

Gettier cases raise some additional problems that are specific to pragmatic accounts of knowledge. The subject's practical situation is irrelevant to our epistemic evaluations in Gettier cases. When a subject's true belief is Gettiered, she fails to know no matter what practical situation she is in. The irrelevance of practical factors in these cases pose a problem. As previously noted, pragmatic accounts typically claim that the practical situation of the subject determines how strong of an epistemic state the subject must be in if the subject is to be in a position to know. And so they are relevant when we determine whether or not a true belief counts as knowledge.¹¹ However, in Gettier cases, the subject's epistemic state is not strong enough regardless of what practical situation the subject is in. So our Gettier intuitions appear to undermine this particular pragmatic claim about what it is to be in a position to know.

¹⁰While the knowledge-first epistemologist may be happy with this type of analysis, I shall not be discussing this approach.

¹¹If the belief is false or the subject lacks a belief, then pragmatic factors are obviously irrelevant. However, as noted above, the pragmatic view concerns the difference between true belief and knowledge, which is accounted for in terms of whether or not the subject is in a position to know.

The second and more important problem arises from the fact that Gettier cases have taught us that knowledge is incompatible with epistemic luck. And as we have learned, the most plausible anti-luck epistemologies are externalist epistemologies.¹² In order to identify when one is in a position to know, one must account for factors that are external to the subject's internal state. The turn to reason-theoretic analyses could be motivated by the need to provide a pragmatic and externalist account of knowledge.¹³ After all, on some views, what counts as a reason is independent of the internal state of the subject.¹⁴ While the appeal to reasons may sometimes be insightful, it does not appear to provide any insight into Gettier cases. As we have seen, reason-theoretic accounts either fall prey to Gettier-style counterexamples or fail to provide accounts with sufficient explanatory power. Unless we have an independent account of the type of reason that knowledge is or provides, the appeal to reasons will be unhelpful.¹⁵ I do not mean to conclude that the reason-theoretic approach cannot work. Rather, I am simply raising an internal challenge for those who adopt the pragmatic viewpoint.¹⁶ We need some explanation of why Gettiered beliefs do not count as knowledge.¹⁷

The decision-theoretic approach offers an alternative but faces serious problems. Decision theory is typically understood as offering coherentist accounts of rational deliberation and on these accounts, the reasonableness of one's deliberations depends solely on the internal state of the decision maker. So decision theory seems unsuitable as a basis for developing an externalist epistemology. As a result, for pragmatists who are partial to the decision-theoretic apparatus, the challenge is to show how we can develop an anti-luck epistemology within this framework. This is the challenge I hope to address.

In summary, Gettier cases pose two specific challenges for the pragmatic account:

1. We must explain why the practical situation of the subject typically determines the standards

¹²Adopting an infallibilist internalist epistemology, one may be able to rule out luck. However, such a view tends to result in skepticism. See (Heller 1999) and (Pritchard 2007a) for further discussion.

¹³From personal discussions, I have found that the proponents of the reason-theoretic approach are not, in fact, motivated by these concerns. Instead, I believe that the use of a reason-theoretic framework allows one to focus on the relevant aspect of reasonable deliberation without committing oneself to any particular theory of rationality.

¹⁴(Finlay and Schroeder 2015)

¹⁵A detailed pragmatic account of knowledge appears to require a detailed account of deliberation. While I am not in principle against the reason-theoretic approach, I am skeptical that such an approach will be fruitful. Reason-theoretic accounts of deliberation tend to be more coarse-grained than their decision-theoretic counterparts. For example, it is hard to talk precisely about the strength of one's reasons in the reason-theoretic framework. However, this simply speaks to the need for richer reason-theoretic accounts.

¹⁶In (Kim 2017), I lay out a variety of problems for the pragmatic viewpoint, identifying some as internal to the project and others that are external (i.e. attempt to criticize and undermine the project).

¹⁷Even for those who reject the pragmatic view, there is no agreement as to what the right explanation is. However, the problem with current pragmatic analyses is that they do not appear to offer *any* explanation.

required for one's true belief to count as knowledge yet, at the same time, is irrelevant for the assessment of knowledge in Gettier cases.

2. We must show how it is possible to develop an illuminating account of knowledge that both accounts for its practical role and rules out epistemic luck.¹⁸

On the view I will develop, the practical situation of the subject does typically determine the epistemic standards that must be met for true belief to count as knowledge. However, I will also argue that there is a minimum standard that must be met in any given practical situation and in Gettier cases, the subject fails to meet this minimum standard. As a result, the subject fails to be in a position to know no matter what practical situation she is in. As it turns out, in order to articulate an appropriate minimum standard for knowledge, I will need to develop an externalist version of decision theory. The decision theory will take an externalist turn by allowing for external, objective factors to come into play when framing decision problems. This externalist move is motivated by the fact the epistemic evaluation of knowers arises from a third-person, objective point of view. So the externalist revision is required in order for decision theory to provide the appropriate evaluations. Therefore, I will address the first challenge by first addressing the second challenge. On the resulting view, a true belief counts as knowledge just in case, from the externalist, decision-theoretic point of view, the subject's total evidence favors the target proposition and the belief is suitable for rational decision-making.

2 An Externalist Decision-Theoretic Epistemology

In response to Gettier cases, many have amended the justified true belief account of knowledge by replacing the justification condition. I shall adopt this approach by using "warrant" to serve as a placeholder for the new condition, which is sufficient to turn true belief into knowledge.¹⁹ However, given our pragmatic aims, we must replace the purely truth-relevant justification condition with one that accounts for practical factors. So let us use the notion of *pragmatic warrant* as the placeholder for whatever it is, from the pragmatic point of view, that turns true belief into knowledge.

¹⁸In the discussion that follows, I'll argue that Gettier cases are at an interesting point in the spectrum of cases where epistemic luck is involved. And the general framework that I develop leaves open the question of how much epistemic luck is incompatible with knowledge.

¹⁹Warrant is the "elusive quality or quantity enough of which, together with true belief, is sufficient for knowledge." (v, Plantinga 1993)

In order to identify the conditions under which a belief is pragmatically warranted, we need to identify the conditions under which it would be appropriate to believe given its practical role. So we should first specify the practical role of belief. Pragmatists have focused on the fact that when you believe that P, you take P for granted when deliberating and deciding what to do.²⁰ When we engage in belief-based reasoning, we engage in coarse-grained reasoning on the basis of the truth or falsity of propositions without taking into account uncertainties. In general, categorical attitudes play an important role in the lives of bounded agents with limited resources.²¹ Beliefs, in particular, simplify one’s deliberations by ignoring likelihood and focusing on truth or falsity. So in order to identify the conditions under which a belief is pragmatically warranted, we need to identify the conditions under which it would be reasonable to take a proposition as true in one’s deliberation rather than consider the likelihood that a proposition is true.²² One intuitive proposal is that such a simplification is warranted only if it makes no practical difference to one’s deliberations.²³ And such a simplification will make not a practical difference just in case the simplifying effect of that belief does not undermine the rationality of deliberation.

In order to identify when the simplifying effect of a belief does not undermine the rationality of deliberation, we will first need some account of what it means to deliberate in a rational way. For such an account, let’s turn to the Bayesian view of rational choice.

2.1 Decision Theory: Part I. Bayesian decision theory (henceforth BDT) articulates what it is to have a coherent standard of evaluation (i.e. a coherent state of mind) for a decision problem. The Bayesian accomplishes this task by first specifying what a decision problem is and identifying the judgments that are relevant for evaluating the choices that one is deliberating between. By then positing a set of coherence constraints governing these judgments, the Bayesian articulates what it is to have a coherent standard of evaluation. And for the Bayesian, a decision maker (henceforth the DM) rationally deliberates just in

²⁰ (Ganson 2008) defends a view of this kind. However, at the moment there has not been any detailed discussion of what it means to take p for granted in one’s practical reasoning. The discussion in the literature assumes that different descriptions of this activity are equivalent for the purposes of the discussion. Some proposals include using p as a premise in reasoning, acting on p, and simplifying one’s decision by assuming that p. It’s not entirely clear if these are equivalent descriptions. I explore some potential differences in section 4.

²¹See (Thomason 1986), (Thomason 2007), and (Thomason 2014).

²²(Locke 2015) offers a similar discussion but focuses on the notion of premising rather than the notion of belief. The discussion there also differs in that the focus is on what is rationally permissible or appropriate rather than on what is warranted.

²³This necessary condition has been helpfully called “practical adequacy” by Anderson and Hawthorne forthcoming, but it is important to note that this is only a necessary condition for pragmatic warrant. In order to provide necessary and sufficient conditions, we need to account for the externalist component of warrant.

case she coherently evaluates the choices that are available in a given decision problem.

To introduce the Bayesian view of a decision problem, consider a case where Catherine is deciding how to get home from the museum. The Bayesian assumes that the decision maker (henceforth the DM) identifies a set of *acts* that she is capable of performing and considers worthwhile in evaluating. Catherine considers two acts, taking the train or a cab. Next, it is assumed that the DM has identified the *consequences* these acts will have for each member of a set of mutually exclusive *states*. Here, ‘states’ refer to possible states of the world. For simplicity, let’s start with a very coarse-grained set of states and consequences summarized in the following decision table (Table 1):

	Traffic	No Traffic
<i>Train</i>	Long Trip, Low Cost	Long Trip, Low Cost
<i>Cab</i>	Long Trip, High Cost	Short Trip, High Cost

Table 1: Catherine’s Decision Problem

In assessing these acts, the Bayesian assumes that the DM’s evaluation of her choices is determined by two independent factors, her desire for the consequences and their likelihood. Thus, Catherine’s evaluation of the acts depends only on her desire for the consequences and her belief about about the states.²⁴

Given this view of decision making, the Bayesian identifies three types of judgments that are relevant for evaluating a set of acts. Call these *deliberative judgments*. First, the DM must assess the desirability of the consequences using her *deliberative desires*. Next, the DM must assess the likelihood of each state using her *deliberative beliefs*. Finally, the DM must compare acts and specify a set of *deliberative preferences* between them.

The Bayesian offers two logically equivalent ways of articulating the coherence constraints that govern these deliberative judgments. The first way articulates rational constraints on deliberative preferences.

One example of such a constraint is that rational preferences are transitive. If a DM prefers act A to B

²⁴Alternatively, we could consider Catherine’s belief that there is traffic given a particular act. However, for ease of discussion, I will work within a theory like the one presented in (Savage 1972) and only consider examples where both act-state independence and the determinism of consequences holds. While many of the restrictions of Savage’s theory are not important to my discussion, I avoid appealing to a theory like the one presented in (Jeffrey 1983) because, to my mind, an account on which belief and desire are separable is more compatible with the constructive approach that shall be offered below.

and B to C, then the DM should also prefer act A to C. The second way is to identify a class of models that represent every coherent set of deliberative judgments the DM could have about any decision problem that is suitably described. The model-theoretic approach claims that any coherent set of judgments must be representable by the type of model proposed. The Bayesian argues that deliberative beliefs, desires, and preferences are coherent if and only if they are respectively representable by probabilities, utilities, and probability weighted utilities. The probability and utility models function as a regulative ideal for our deliberative judgments since they represent the type of rationally permissible judgments that a DM can have about her decision problem. These two ways of articulating coherence are logically equivalent because it can be shown that the DM has a coherent set of preferences if and only if the DM's deliberative beliefs, desires, and preferences can be represented by probabilities, utilities, and probability weighted utilities.²⁵

2.2 Pragmatically Warranted Belief. Now that we have summarized our theory of rational deliberation, we can begin to answer the question with which we started. When is a subject's belief pragmatically warranted? When may a subject take a proposition for granted in her deliberation? Since the standard decision-theoretic framework only deals with graded notions like probability and utility, there is no explicit place for the all-or-nothing beliefs that are typically understood to be necessary for knowledge. However, a natural account can be provided.

Within a well-defined decision problem, a set of coherent preferences is enough to entail that one's beliefs, desires, and preferences are representable as probabilities, utilities, and probability weighted utilities.²⁶ Call such coherent judgments *ideal* and the set of these coherent judgments a *rational ideal*. Rational ideals function as a rational standard for decision problems.²⁷ And since the purpose of engaging in practical deliberation is to make a choice, the primary function of a rational ideal is to identify which choices count as rational for a given decision problem. On the Bayesian account, a choice is rational just in case it is amongst the most preferred (i.e., maximizes expected utility). Our intuitive proposal was that a belief is pragmatically warranted only if it makes no difference to one's practical evaluations.

²⁵(Savage 1972) offers the paradigmatic proof.

²⁶While coherent preferences for any particular problem cannot ensure a unique representation, this will not matter for our purposes. While I will talk about conditional expected utilities in addition to talk about conditional preferences, we could simply use the latter way of speaking. However, for ease of presentation, I will appeal to both locutions.

²⁷There's an interesting question whether we ought to be coherent across decision problems and if we should, whether this results in a demand to be coherent across every decision problem. I have discussed these issues in (Kim 2014).

If we situate this proposal within the Bayesian view of rational deliberation, we conclude that for a given decision problem, a DM’s belief is pragmatically warranted only if her preferences given this belief are *practically coherent* with her ideal preferences. Two sets of preferences are practically coherent just in case they identify the same set of choices as rational. Thus, we have one initial condition governing pragmatically warranted belief: S’s belief that P is pragmatically warranted only if S’s preferences given her categorical belief that P are practically coherent with S’s ideal preferences.²⁸

This principle proposes that the reasonableness of an all-or-nothing belief depends upon the resulting stability in the evaluation of one’s choices. Pragmatically warranted beliefs do not alter what counts as rational relative to our more fine-grained beliefs. To illustrate, suppose Catherine has the following probabilities and utilities (Table 2):

	Traffic [25%]	No Traffic [75%]
<i>Train</i>	Long Trip, Low Cost (.5)	Long Trip, Low Cost (.5)
<i>Cab</i>	Long trip, High Cost (.2)	Short Trip, High Cost (.7)

Table 2: Catherine’s Decision Problem with [Probabilities] and (Utilities)

Calculating expected utilities, $EU(Train) = .5$ and $EU(Cab) = .575$ so taking a cab is ideally preferable to taking the train. Calculating expected utilities conditional on $[\neg T]$ there being no traffic, $EU(Train|\neg T) = .5$ and $EU(Cab|\neg T) = .7$. So even if she believes that there is no traffic, she still prefers the cab. The resulting stability of what Catherine prefers underwrites the rationality of her belief that $\neg T$. Of course, given different probabilities or utilities, she might not be rational taking $\neg T$ for granted. Keeping all else fixed, if a short, high cost trip had a utility of less than .6, $EU(Train) > EU(Cab)$ but $EU(Train|\neg T) < EU(Cab|\neg T)$. Her belief that $\neg T$ would also be unstable if we kept the utilities fixed and $p(T) > .4$.

An interesting problem arises in cases of dominance, where the consequences of one choice are always more desirable than the consequences of all the others. No matter what the DM takes for granted, she will

²⁸Sets of preference can be practically coherent without producing the same preference ordering. One’s unconditional preference ordering may be $A \succ B \succ C$ and one’s preference ordering given P may be $A \succ C \succ B$. Thus, there is a question about whether pragmatically warranted belief must also require a stable preference ordering. Since the difference between these two views will not make a difference to my discussion of Gettier cases, I will not consider this issue in any detail.

always prefer the same choice. The DM would thereby be rational in taking any proposition for granted. However, it is unintuitive to think there are situations in which one may arbitrarily either believe P or believe $\neg P$ and be rational in doing so. Believing that P commits the subject to the truth of P . Moreover, truth is an aim and norm of belief.²⁹ So one ought to have some epistemic reason for believing. We can ensure that the DM's beliefs are not epistemically arbitrary by adding in the minimal condition that the DM must, on the whole, possess evidence that favors P over $\neg P$. So when we evaluate the DM's epistemic state, their total evidence must favor P . While this is a somewhat vague condition, it will do for the moment; I will offer more discussion of this minimal condition in section 3. Thus,

*S's belief that P is pragmatically warranted only if S's preferences given S's belief that P are practically coherent with S's ideal preferences and S's total evidence favors P over $\neg P$.*³⁰

This latter necessary condition will serve as the minimum epistemic standard for knowledge and failing to meet this standard would mean that one's belief is not epistemically motivated. If one's belief fails to be epistemically motivated, then one fails to know in every practical situation. In section 3, I will show that, on the proposed account, Gettier cases lie at a very specific point on the spectrum of epistemically unmotivated beliefs. However, the condition as it has been described so far will not do. After all, subjects do appear to possess evidence that their Gettiered beliefs are true and so they satisfy this minimal condition. What we need is a somewhat specific measure of a subject's strength of evidence. And to develop this, I will outline a decision-theoretic framework that will allow us to measure one's strength of evidence in a way that depends upon what counts as evidence. And by fixing what counts as evidence from an external point of view, we will be able to consider external, "objective" measures of evidential strength when evaluating DM's. And we shall find that from this perspective, Gettiered subjects possess no evidence for or against the target proposition.³¹

²⁹There is a lively debate about what it means for belief to aim at the truth. I do not believe that the point here depends upon any particular interpretation of this platitude. See (Chan 2014) for a thorough discussion of the topic.

³⁰Though (Fantl and McGrath 2002) presents a necessary condition for rational belief that is similar, I will develop its connection with knowledge in a way that is very different from their account. In addition, we should read the condition that one must have favoring evidence as a necessary condition for a subject to possess propositional justification. So if I am being careful, I am only proposing an account of when one is in a position to know. To turn such an account into an account of knowledge, one must have a satisfactory account of doxastic justification as well. Finally, on my account, all-or-nothing belief is independently necessary for knowledge. I reject the reduction of all-or-nothing belief to degrees of belief that possess some property. Problems for reductive accounts of rational belief are canvassed in (Ross and Schroeder 2012). By requiring that the subject possess a categorical belief, we can avoid some of these problems.

³¹ Irrelevant possibilities, which are those whose truth or falsity have no relevance for the outcomes of a choice, also pose a problem for the account. After all, if we have favoring evidence for an irrelevant possibility, the account appears to entail that we know that such a possibility obtains. I address this issue in fn. 43.

2.3 Decision Theory: Part II. So far, our analysis of pragmatically warranted belief has focused on coherent belief. We have proposed that a pragmatically warranted belief must be practically coherent and epistemically motivated. However, if we are to identify the conditions that are both necessary and sufficient for pragmatic warrant, we also need to ensure that pragmatically warranted belief rules out epistemic luck. So pragmatic warrant cannot simply be a matter of internal coherence. The problem is that the Bayesian account of rational deliberation is only concerned with coherence. So let's return to the Bayesian account to see how we might broaden its scope.

Decision theories answer three questions when explicating what it means to have a rational standard of evaluation.

1. What is a decision problem?
2. Which types of judgments are relevant for evaluating choices?
3. What constraints must these judgments meet in order to count as rational?

Decision theorists typically focus on the third question, and there has been a long debate about various coherence constraints on beliefs, desires, and preferences.³² But how does the Bayesian answer the first two questions? The Bayesian proposes that decision problems can be described by a set of acts, states, and consequences. And the deliberative preferences, beliefs, and desires over these acts, states, and consequences are all that is relevant for evaluating one's choices. Unfortunately, these answers leave us with an incomplete account of rational deliberation. Similarly, suppose one proposed a purely descriptive account of decision making that focused solely on the evaluative component of choice. This account would be problematic since it would miss out on the all important framing component of choice.³³

2.3.1 Two Limitations of Bayesian Decision Theory. The first problem is that the Bayesian offers no detailed account of how to frame a decision problem.³⁴ It is simply assumed that there is some way

³²See (Gilboa 2009) for a summary of the debate.

³³(Tversky and Kahneman 1981), (Kahneman and Tversky 1984), and (Tversky and Kahneman 1986)

³⁴There are, of course, some constraints on what counts as an appropriate framing. For example, the principle of act-state independence is a constraint on how a decision maker models his or her decision problems. However, most decision theorists think that we do not need to provide a comprehensive answer to this question. For example Savage writes, "I believe, and examples have confirmed, that decision situations can be usefully structured in terms of consequences, states, and acts in such a way that the postulates of [Foundations of Statistics] are satisfied. Just how to do that seems to be an art for which I can give no prescription and for which it is perhaps unreasonable to expect one." (79, Drèze 1990)

that we can do this. This assumption makes sense when one is choosing between well-defined bets. If I offer you a choice between two bets, one offering \$7 if a coin lands heads and nothing otherwise, the other offering \$10 if the same coin lands tails and nothing otherwise, the decision table is fixed. The relevant states, which represent possible answers to the question of which side the coin will land, are those that matter for determining the outcome of the bet. The relevant consequences are the monetary prizes. However, when Catherine decides between taking the train or a cab, it is far from obvious what the relevant states and consequences ought to be. When evaluating Catherine's attitudes about her decision problem, how should we frame her decision problem? Should we incorporate into her decision problem considerations about the weather? What about the possibility that the train may derail or that the cab may crash? In order to have a comprehensive account of rational deliberation, we need to answer the framing question: how should the DM's choice problem be framed?

Second, the Bayesian solely focuses on deliberative judgments – those judgments that concern the acts, states, and consequences of a decision problem – and assumes that such judgments are all that matter for characterizing rational deliberation. Viewed as a comprehensive theory, BDT states that so long as a DM deliberates on the basis of any coherent set of beliefs, desires, and preferences, she is free from any rational criticism. This claim to comprehensiveness would only be correct if our deliberative judgments were not based upon any other type of judgment or attitude. However, over the last forty years, an overwhelming amount of research, beginning with the preference reversal experiments of Slovic and Lichtenstein, has slowly undermined this assumption.³⁵ These experiments show that many of the judgments we use to deliberate are deeply context-dependent, changing from one deliberative context to the next. In reply, many researchers have adopted the methodological assumption that our deliberative judgments are *constructed*. Slovic and Lichtenstein write that “the big picture is the overwhelming evidence that people often do not have preexisting preferences but must construct them to fit the situation.”³⁶ If we adopt this constructive point of view, a comprehensive theory of rational deliberation must also account for how our deliberative judgments are constructed in each context. So BDT is incomplete.³⁷

³⁵See the articles in section 2 of (Lichtenstein and Slovic 2006).

³⁶(Lichtenstein and Slovic 2006), p.12

³⁷In (Kim 2014), I offered a detailed discussion of the challenges raised by choice behavior research for normative decision theory. I also show the constructive framework described here can address these challenges. As I note there, the proposed framework has been heavily influenced by (Raiffa and Keeney 1976).

In order to address these two limitations, we must add to the Bayesian theory some account of how to frame a decision problem and some account of how DMs construct their deliberative judgments. I will call the resulting theory, a **Constructive Decision Theory**.³⁸ I call the view constructive to make it explicit that decision problems are not always given and that deliberative judgments do not always pre-exist. Both must sometimes be “constructed”.

2.3.2 Constructive Decision Theory. In order to articulate how to frame a decision problem (i.e., demarcate the set of relevant states and consequences), we are guided by the following motto: *what matters in deliberation depends upon what deliberation is for*. And I shall begin our discussion by considering how this might be done from the first-person point of view. We shall later broaden the scope of our inquiry to consider how objective considerations might factor in when we frame the DM’s decision problem from the third-person, evaluative perspective.

If we view table 2 as Catherine’s own view of her decision problem, then the decision matrix indicates that she restricted the set of relevant consequences to only those that described the time and cost it would take to get home. This is all she took into account and all she cared about. So it seems that only two values were relevant to her deliberation, expediency and thrift. By specifying the relevant values, one can determine the set of relevant consequences. Though it may be a bit strange to suppose that Catherine explicitly judges that these two values are all that matter, it is quite natural to state that she has the goal of getting home in the fastest, most cost-effective way. Goal judgments offer an intuitive way in which the DM demarcates a set of relevant values. Once a goal is set and a set of values is selected as all that matters, the relevant consequences are simply those that describe the outcomes of one’s actions relative to these values. So Catherine only needs to describe those consequences of her actions that matter given her concern for the duration and cost of transportation. As a general rule, *the consequences that count as relevant depend upon the values that count as relevant*. What is notable about our goals is that they may change from one deliberative context to the next. In another context, Catherine may have the goal of getting home in the least stressful, most aesthetically pleasing way. This in turn will identify a different

³⁸Some Bayesians would dismiss such judgments and argue that for a normative decision theory, every state and consequence is taken into account and our attitudes are ideally captured by a single probability and utility function (c.f. Jeffrey 1965) On this view, the only reason we restrict the possibilities we consider comes from the economic costs of taking everything into account. However, for some non-economic reasons for restricting the possibilities that we consider, see (Shafer 1986).

set of relevant values. Thus, what deliberation is for is deliberation-specific so the consequences that count as relevant are also deliberation-specific.

Catherine must also demarcate a set of relevant states. I proposed that the only question she considered in her deliberation was whether or not there would be traffic on the streets. Neither the color of the cab, the cleanliness of the train, nor the personality of the cab driver were deemed relevant. Since we have a rule for demarcating the set of relevant consequences, we also need one for our states. What matters must depend upon what deliberation is for and the relevant states are indirectly determined by our goals. Some states matter because their actualization will affect whether one or another of the relevant consequences occur. Catherine recognizes that if there is traffic on the streets, the cab ride may be costly and slow. Since she cares about the cost and duration of the trip, she should consider the possibility of traffic. So the first general rule is that the DM must consider all the states that she thinks will affect whether one or another of the relevant consequences occurs. Call these *consequence-determining states*.

A second set of possibilities matters since their actualization is relevant for assessing the likelihood of the consequence-determining states. Suppose Catherine would find it very likely that there would be traffic if the President were visiting the city. If she judges that the President's presence is relevant for assessing the likelihood of a consequence-determining possibility, Catherine should consider this possibility as well. Call these *evidentially-relevant states*.

So one possibility is that the DM herself makes certain judgments, such as goal judgments, that set the parameters for her decision problem. Call these *constructive judgments*, which allow the DM to demarcate a set of relevant consequences and states. These same judgments can also help to construct one's deliberative judgments. They do so by placing constraints on the set of rationally permissible deliberative judgments. For example, the set of relevant values places constraints on the utilities that are rationally permissible. If the cost and duration of the trip are the only values that matter to a deliberation, then a short, low cost trip should be more desirable than a long, high cost trip since the latter outcome is better according to both values.³⁹

The type of constructive decision theory that I have so far described is compatible with the Bayesian

³⁹This merely induces a partial ordering but if we also assume that goals have more structure then we can get more uniqueness. For example, in the extreme case where a goal determines the relative importance of values then we may be left with a utility unique up to positive linear transformation.

account. Rational deliberative beliefs, desires, and preferences are still probabilities, utilities, and probability weighted utilities. The only change is that we incorporate how a DM might set the parameters of her deliberation by accounting for constructive judgments within our theory. Though this expands the scope of the types of judgments that are within the purview of our decision theory, it also narrows the scope of the Bayesian coherence constraints.

On one interpretation of BDT, the entirety of one's beliefs, desires, and preferences ought to be representable by a single probability, utility, and probability weighted utility function. So long as there is no change in what the DM believes or desires, these probabilities and utilities should then be used for every decision problem. This interpretation is unrealistic and incomplete since it fails to account for the context-sensitive construction of our deliberative judgments. The constructive account described above proposes that our deliberative judgments are often constructed on the fly and may vary from one deliberative context to the next. If constructive judgments are deliberation-specific and deliberative judgments are constructed on their basis, the beliefs and desires that count as rational for one deliberation are not necessarily what counts as rational for another. Since the Bayesian principles of rationality explicate what it is for these deliberative judgments to be coherent, they must only apply within a particular deliberative context. So the demand that deliberative beliefs, desires, and preferences be representable by a probability, utility, and probability weighted utility is restricted to the decision problem for which these judgments are constructed.⁴⁰

2.3.3 Beyond coherence. While the constructive decision theory described in the previous section remains a coherentist (and internalist) account of rational deliberation, the recognition that every comprehensive decision theory must account for how to frame a decision problem opens up the decision-theoretic framework to the development of externalist standards. To do so, we will first show how external, objective factors can be appealed to in setting the parameters of a decision problem. More importantly, these objective factors are relevant when we evaluate DMs from the third-person perspective. And this evaluative use of decision theory is what's relevant for a theory of knowledge. Thus, by appealing to externally-framed decision problems, we will be able to develop an externalist, decision-theoretic epistemology.

⁴⁰In (Kim 2014), I explore how this very limited interpretation of BDT could be used to account for the more global aspects of instrumental rationality.

In order to deal with Gettier cases, we will only need to reconsider how to determine which states are evidentially-relevant. When we pick out a set of evidentially-relevant states from the third-person perspective, we are demarcating what counts as evidence and how much that evidence counts in the context at hand. This notion of “counting as evidence” is crucial if we, as bounded agents, are to have clear and meaningful measures of evidential support.

For example, suppose we consider a two-sided coin and try to assess the likelihood that it will land heads on the next toss. We know that this coin has been tossed a thousand times, and we can ask anything about these tosses. What sort of information is relevant when assessing the likelihood that coin will land heads on the next toss? This question is relevant if we are to evaluate the strength of a subject’s evidence, when the subject possesses information about these previous tosses. The answer to this question depends upon what is taken into account, what information counts as evidence. And what counts as evidence depends upon certain assumptions that must be made for inquiry (or for our purpose, the evaluation of a subject’s strength of evidence) to proceed. On the assumption that the coin is fair, no information about the previous tosses counts as evidence for or against. Alternatively, on the assumption that the coin could be biased, information about the proportion of heads in the previous tosses would be relevant. But is the order of the outcomes relevant information? This depends upon whether or not it is assumed that the coin flips are independent. What about the color of the coin? Such information would count as relevant within a context where we consider the possibility that the coin’s color could have an effect on its bias. As one can see, such questions go on *ad infinitum*. Thus, we must restrict what can count as evidence in order to engage in inquiry and to evaluate strength of evidence.

Both subjective factors and objective factors can be used to determine what is evidentially-relevant. Suppose Catherine is trying to decide whether to take a cab or the train back home from the museum. And suppose that the only pair of consequence-determining states, from Catherine’s points of view, are the ones in which there is or is not traffic on the road. One way to determine the states that are relevant for assessing the probability of these consequence-determining states is to consider Catherine’s actual judgments (i.e. what Catherine actually considers). For example, consider the case in which Catherine takes into account the possibility that the President is visiting the city and also judges that this possibility is relevant for assessing the likelihood of traffic. Such a possibility is evidentially-relevant and since the

explanation of its relevance only appeals to the DM's own judgments about what counts as relevant, it is relevant from the subject's own point of view. So evidentially-relevant states can be demarcated in terms of what is subjectively relevant.

Alternatively, we can demarcate what is evidentially relevant from a more objective point of view. For example, consider a case in which there has been an accident on the streets. Catherine is unaware of this fact and has not considered this possibility. Thus, she has not judged the possibility of an accident as being relevant for her deliberation. So from her point of view, information about accidents on the streets are ignored and do not count as evidence. However, there is an intuitive sense in which this possibility ought to be considered. After all, if Catherine had taken it for granted that there had been a car accident, she would have come to a very different assessment of the likelihood of traffic. The intuitive notion of objective relevance comes from the observation that true but unconsidered propositions relevant for assessing the likelihood of any relevant state ought to be taken into account. More on how to interpret "likelihood" in a moment.

To put the point another way, from the third-person point of view, facts about the actual situation determine the space of evidentially-relevant possibilities within which one should ideally inquire. And it is relative to this space that we evaluate the strength of a subject's evidence. After all, suppose that Catherine thinks it unlikely that there is traffic since it is fairly late in the evening. Thus, from her point of view, she possesses strong evidence. In contrast, from our perspective, her evidence is much weaker. One way to account for this difference is to incorporate these objectively-relevant possibilities into the decision space.⁴¹ And from this new point of view, Catherine's evidence no longer justifies a high degree of confidence.

There are two ways of explaining why, from the third-person perspective, Catherine's strength of evidence is lower once the possibility of an accident is incorporated into the space of possibilities. First, we could counterfactually consider Catherine's rational degree of belief within this broader space. Since Catherine's evidence cannot discriminate between there being an accident or not and since there will almost certainly be traffic if there were an accident, she should be less confident when constructing her

⁴¹The set of states is a boolean algebra so if the possibility of an accident is incorporated then so much the possibility that there is no accident on the streets.

probabilities in this broader space of possibilities. Alternatively, we might propose that there is some evidential probability that there is traffic given Catherine's evidence.

I am agnostic between these alternative ways of understanding how changes in the space of possibilities can alter subjective or objective probabilities. Thus, I do not have a theory or procedure to identify the set of relevant states from the objective point of view. I am merely relying upon the intuitive idea that when we evaluate a DM's evidence, we do so from the point of view where many true but unconsidered propositions relevant for assessing the likelihood of any relevant state are taken into account. Moreover, these possibilities alter the probability by altering what counts as evidence and how much the evidence counts. While I shall rely upon this intuitive characterization, I have included some more discussion of objective relevance in Appendix A to address some questions and problems.⁴²

2.4 The Pragmatic Account of Knowledge. We have now seen how both subjective and objective factors can be used to set the parameters of a decision problem. So we can now differentiate various framings of decision problems in terms of the types of parameters they take into account. It should be noted that the use of subjective or objective decision parameters is not mutually exclusive. There are framings that take into account both objective and subjective factors, and there are framings that take into account just one or the other. By adopting the constructive decision-theoretic framework and considering framings of decision problems that incorporate objective parameters, the internalistic view of rational deliberation presented by BDT can be transformed into a view that is friendly to the externalist. And this externalist decision-theoretic framework can be used to evaluate the strength of a subject's evidence from the third-person point of view. As I noted, this more objective measure of a DM's strength of evidence can either be understood as evidential probabilities or the DM's rational degrees of belief relative to the objectively-defined set of states. We can thereby appeal to objectively-framed decision problems to produce an account of knowledge as pragmatically warranted belief.

Pragmatically Warranted Belief: S's belief that P is pragmatically warranted if and only if, in an objectively-framed decision problem, S's preferences given S's belief that P are practically coherent with S's ideal preferences and S's evidence favors P .⁴³

⁴²My discussion will focus on two problems. The first problem is that almost everything can be relevant. The second problem is that we can gerrymander possibilities so that they are relevant.

⁴³For our purposes, the decision problem counts as objectively-framed if the set of evidentially-relevant states accounts for evidentially-relevant matters of fact. However, there may be reasons to incorporate other objective parameters into the analysis of pragmatically warranted belief. While it will go beyond the scope of this paper to discuss other objective decision

3 Gettier Cases

The pragmatic account imposes a minimum standard for knowledge whereby the subject must have favoring evidence from the evaluative, third-person point of view (i.e. relative to objectively-framed decision problems). I previously noted that this minimum standard for knowledge would not be met when a subject's belief is Gettiered, and when this minimum standard is not met then the subject fails to know no matter what practical situation she is in. To show this, I will compare and contrast skeptical possibilities and Gettier possibilities. By doing so, we will be able to get a clearer understanding of what it means to evaluate the strength of a subject's epistemic state relative to one or another way of fixing what counts as evidence and how much that evidence counts.

Under normal circumstances, skeptical possibilities, such as brain-in-a-vat possibilities, are ignored. However, it's easy to enter into a context (e.g. a philosophical discussion) where such possibilities are taken seriously. We may do so for the sake of engaging in a certain epistemological project. As Barry Stroud has noted, radical skepticism is closely connected to the epistemological project of trying to ascertain whether we can know anything about a particular domain without presuming any knowledge of that domain.⁴⁴ Of course, when this possibility is taken seriously, then we cannot assume to possess any evidence for or against any proposition about the external world. For example, perceptual experiences that would normally ground our judgments about having hands no longer count as evidence in favor of such claims.

Thus, in a context where BIV possibilities are deemed serious and relevant, what counts as evidence for

parameters in any detail, it may be useful to consider another example. Goal judgments are subjective parameters that can be used to identify the relevant consequences. However, if there are objective values or goals that a subject ought to have, then these can be used as objective parameters that demarcate the relevant consequences. These alternative parameters may be relevant for the pragmatic account. Some have criticized pragmatic accounts as presenting a highly unstable account of knowledge. For example, (Anderson and Hawthorne forthcoming) present cases where the option to double-check alters one's epistemic status. We can prevent this type of instability if the set of available actions is fixed from the objective point of view. Either the double-checking option is available or it is not and so considering such an option as available does not change whether or not one knows. If we embrace the full set of objective parameters, there will be a way of framing decision problems such that in the actual world, when we fix facts about both the present and the future, there is a unique decision problem the agent faces. This move, however, would not remove instability across possible worlds. I personally think that there just has not been enough discussion about how stable knowledge actually is. For example, the paradoxes present in (Kripke 2011) raise fascinating questions about the stability of knowledge. And so I am unsure as to how to evaluate these objections about the instability of knowledge on the pragmatic account. In addition, I raised the problem of irrelevant possibilities in footnote 31 and can appeal to the notion of an objectively-framed decision problem to offer a tentative response. Objectively-framed decision problems only include relevant possibilities. Since the theory only explains how to evaluate deliberative beliefs (i.e. beliefs within a deliberative context), it is silent on whether we possess knowledge of irrelevant possibilities. While I'm somewhat agnostic about how to fill in this silence, I am partial to a Peircean response. When we attribute knowledge that P to ourselves or others, we are interested in whether the question, "P or not-P?", has been adequately resolved by the subject. Within any particular deliberative context, questions about irrelevant possibilities are what Peirce would call fake "paper" questions. And if no question has really been asked, who cares what the supposed answer is.

⁴⁴(Stroud 2000)

or against propositions about the external world changes. And from this perspective where skeptical possibilities are taken seriously, we lack favoring evidence for external world propositions.

In this respect, Gettier possibilities function exactly like skeptical possibilities. Consider Chisholm's case in which you look out at a field and see what looks exactly like a sheep and come to believe that there is a sheep in the field. Unbeknownst to you, the animal you are looking at is a dog, disguised as a sheep. But there is in fact another sheep in the field, behind the hill out of your field of vision. Under normal circumstances, the disguised dog possibility is ignored and not taken seriously. However, suppose that such a possibility were considered. Under those circumstances, your perceptual experience no longer counts as evidence. After all, it is perfectly compatible with the disguised dog possibility. Since the only potential evidence you possess for the claim that there is a sheep in the field is your perceptual experience, you thereby lack favoring evidence for that claim.

Like skeptical possibilities, Gettier possibilities, when taken seriously, undermine one's potential evidence. They do so by altering what counts as evidence in favor of certain propositions. In both cases, once certain alternatives are taken seriously, what had previously counted as favoring evidence no longer does so. There are, however, a few differences between radical skeptical hypotheses and Gettier possibilities. First, skeptical possibilities are more general in that they undermine a whole class of potential favoring evidence. In contrast, Gettier possibilities merely undermine the evidence one actually possesses. After all, one could simply walk up to the disguised dog and walk over the hill to see the sheep. Second and more importantly, Gettier possibilities are actual. The animal is a disguised dog. Jones is lying about owning a Ford. The match has impurities. In contrast, there is no presumption that skeptical hypotheses are true. This difference is crucial if my account is to avoid skepticism.

Constructive decision theory is able to capture the context-sensitivity of what counts as favoring evidence by allowing for the set of evidentially-relevant states to change from one decision context to next. When this framework is used to evaluate whether subjects know, we must ensure that all true evidentially-relevant states are taken into account. This restriction to actualized states is what allows the account to differentiate Gettier possibilities from skeptical ones. From the perspective in which Gettier possibilities are taken into account, subjects lack, on the whole, evidence for the target proposition. We can thereby conclude that when a subject's belief has been Gettiered, she lacks evidence from this

objective, knowledge-evaluating point of view. If she lacks evidence, then she fails to satisfy the minimal epistemic condition required for knowledge. And if she fails to satisfy the minimal condition, then she fails to know in every practical context she could be in.

3.1 The Pragmatic Analysis of Gettier Cases. One may object that I have simply added an externalist and truth-relevant component to the pragmatic account – a component that any good theory of knowledge should have. In addition, it might be thought that the proposed pragmatic account has similarities to both the indefeasibility and relevant alternative analyses of knowledge. So I'd like to make a few remarks about what insight we gain by adopting this pragmatic account. However, before I turn to that task, it is worth reiterating that my primary aim was to address a challenge internal to the pragmatic viewpoint. The challenge was to present develop an externalist account of rational deliberation, whether from the reason-theoretic or decision-theoretic perspective, that could provide any explanation as to why Gettiered beliefs do not count as knowledge. After all, previous accounts could offer no such explanation.

Let us begin by considering the relevant alternatives account of knowledge. This view famously runs into difficulties demarcating the set of alternatives that must be ruled out in order to count as knowing. As is explicit in (Lewis 1996), it is unclear how we can demarcate what counts as a relevant alternative in way that excludes skeptical alternatives but includes Gettier possibilities. The problem in Lewis' account is that it makes an appeal to similarity that is not sufficiently discriminating. The pragmatic notion of objective relevance offers a fairly simple resolution to this problem. First off, the pragmatic account does not talk of ruling out alternatives and instead talks of a space of relevant states or possibilities. Next, the account of relevance is pragmatic in that it ultimately depends upon what is relevant for the decision maker's evaluation of her choices in a decision problem. And I proposed that objectively relevant states are defined in terms of true facts that are evidentially relevant for the assessment of likelihood. As a result, Gettier possibilities count as relevant because they are true while skeptical possibilities do not because they are presumably false.⁴⁵ Thus, my pragmatic account acknowledges the similarity between the two types of possibilities while identifying a small but very important discriminating property.⁴⁶

⁴⁵Of course, we can adopt an evaluative perspective from which skeptical possibilities are taken seriously. We can do so because the subject herself takes these possibilities seriously or we can do so because we as evaluators take these possibilities seriously.

⁴⁶See appendix for more discussion.

The no defeaters account of knowledge runs into difficulties making sense of a defeater. The problem arises from the need to differentiate misleading from defeating evidence. In the well-known case of Tom Grabit and his mother, Tom has stolen a book from the library and Smith has seen him do so.⁴⁷ However, unbeknownst to Smith, Tom's demented mother has asserted to a friend that Tom was nowhere near the library but Tom's twin John, who is a figment of Mrs. Grabit's imagination, was at the library during the theft. This example posed problems for early indefeasibility accounts. The challenge was to provide an account of a defeater such that Mrs. Grabit's assertion was ruled out and dismissed as a misleading piece of evidence. Though subsequent analyses of defeaters were able to rule these cases out, it became increasingly difficult to imagine that any account could demarcate all true facts into those that were genuine defeaters and those that were misleading.⁴⁸

The most obvious and important thing to note is that the pragmatic account makes no appeal to defeaters. In fact, the account has no use for such a distinction. Instead, the account only relies upon a distinction between relevant and irrelevant states so it does not depend upon any distinction between defeating and misleading facts. In fact, all evidentially relevant facts are taken into consideration from the objective point of view. Let me explain. A particular fact is misleading only if it is taken alone (i.e. if the relevant facts are gerrymandered in the right way). In the case of Demented Mrs. Grabit, the fact that Mrs. Grabit told a friend that Tom's twin and not Tom was at the library is misleading only if it is considered independently of the facts that Mrs. Grabit is a compulsive liar and that Tom has no twin brother. The pragmatic account is happy with so-called misleading facts because it proposes to account for all evidentially-relevant possibilities and facts.⁴⁹ And the thought is that once all the objectively relevant possibilities are incorporated, no particular fact can be misleading.⁵⁰

Let us conclude by bringing our discussion up to date. In the recent epistemological literature, the two main competitors for explaining why true belief falls short of knowledge in Gettier cases are the ease of error and lack of credit explanations. "According to *Doxastic* Ease of Error Approach, a Gettiered belief

⁴⁷(Lehrer and Paxson 1969)

⁴⁸See the discussion of defeasibility analyses in (Shope 1983).

⁴⁹Not all seemingly relevant possibilities are evidentially-relevant. Suppose that when Catherine is trying to go home from the museum, a series of incredibly unlikely events prevents an accident. While these events are relevant to the story of why there isn't an accident, they are not relevant to the evaluation of likelihood since $p(\text{Accident} | \text{Unlikely}) = 0$ and $p(\neg \text{Unlikely}) \approx 1$.

⁵⁰The current approach is to be distinguished from the no false lemmas for more obvious reasons. The decision-theoretic account of Gettier cases appeals to a particular measure of the subject's strength of evidence and does not appeal to instances of reasoning. Furthermore, the proposed account works equally well even when there are no relevant false beliefs.

that P is one whose subject S could easily have *believed* falsehoods similar to P , in ways similar to how S actually believes P .”⁵¹ The most natural way to fill in this account is by appealing to counterfactual properties that a subject’s belief must have in order for her true belief to count as knowledge.⁵² According to the lack of credit approach, “ S ’s belief B in P is gettiered iff (B is true and epistemically justified but) B ’s truth isn’t sufficiently creditable to S ’s cognitive abilities.”⁵³ This account has been explicated by appealing to the explanatory salience of the subject’s cognitive abilities as well as the manifestation of the subject’s cognitive abilities.⁵⁴

We can better understand the pragmatic account of knowledge by highlighting its alternative explanation of why true belief falls short of knowledge in Gettier cases. Rather than possessing counterfactual properties or manifesting the subject’s cognitive abilities, a pragmatically warranted belief is one that is well-formed relative to its function in deliberation and the subject’s actual situation. The pragmatic account of knowledge focuses on the activity of believing within the broader activity of deliberating. And rational deliberation manifests more than one’s cognitive abilities. For example, one may have to evaluate outcomes appropriately relative to the values and goals that relevant to a deliberative context. The resulting account proposes that Gettiered beliefs are those that have not been skillfully formed in a deliberative situation.

The pragmatic account also offers a novel view of the significance of Gettier cases. The minimal epistemic condition that I have been working with is very weak and may have some undesirable consequences. I suggest stronger conditions in section 4.2. However, I started with this condition because it brings to light where, from the pragmatic viewpoint, Gettier cases lie on the spectrum of cases. On the proposed view, there may be contexts in which we know somewhat unlikely propositions (e.g. that a fair coin will land heads on one of two tosses). But Gettier cases are special cases where one lacks favoring evidence for either P or not- P . On one side of Gettier cases are those where one irrationally believes that P even though one possesses, on the whole, varying amounts of favoring evidence for not- P . On the other side of the spectrum are cases where one believes that P and possesses, on the whole, varying amounts of favoring evidence for P .

⁵¹p.6, (Coffman forthcoming)

⁵²See (Pritchard 2007b) and (Hawthorne and Lasonen-Aarnio 2009) for some proposals.

⁵³p.11, (Coffman forthcoming)

⁵⁴(Greco 2010), (Sosa 2010)

Since Gettier cases are those in which one possesses neither evidence in favor or against, there is a sense in which, from the third-person point of view, Gettiered belief are true by a matter of pure luck. They are akin to cases in which one thinks that one possesses all sorts of evidence that a coin toss is biased and will land heads. However, as a matter of fact, the coin is fair and so this evidence does not, in fact, favor that outcome. Cases of pure luck are particularly interesting, because they are the paradigmatic cases in which a successful outcome is not, in any sense, due to the agent's skill. Rather, it's solely a matter of luck. These cases mark an important point in the spectrum of cases in which the agent deserves no credit regardless of the practical situation the agent finds herself in.

4 Exploring Pragmatic Epistemologies

At the outset, I proposed to engage in a broad exploration of pragmatic epistemologies. However, I quickly narrowed in on one type of pragmatic proposal. I did so, in part, because I believe that the proposed account would bring to light an interesting feature of Gettier cases, but I also did so because I believed that this was the best way to introduce the constructive decision-theoretic framework, which I find rich enough for exploring pragmatic epistemologies. So to conclude, I'd like to address some objections to the proposed view. By doing so, we will be able to identify various junctures at which alternatives to the current proposal can be explored. My aim here is to show not only that the current framework is a rich and flexible one, but also to show where exploration might be fruitful.

4.1 Belief, practical adequacy and closure. I considered the following necessary condition on pragmatic warrant: S's belief that P is pragmatically warranted only if S's preferences given her categorical belief that P are practically coherent with S's ideal preferences. This type of condition has been called the practical adequacy condition for knowledge.⁵⁵ Call a belief practically adequate just in case it satisfies this necessary condition. Unfortunately, practical adequacy as described is incompatible with certain desirable closure principles on knowledge, such as single-premise closure.⁵⁶ I appealed to this practical adequacy principle because it is intuitive and worked for our purposes, but we can explore variations of the principle that preserve single-premise closure.

Consider the strongest belief (the belief that entails all the others) in the set of practically adequate

⁵⁵(Anderson and Hawthorne forthcoming)

⁵⁶(Hawthorne and Stanley 2008) and (Zweber forthcoming) both discuss this problem.

beliefs.⁵⁷ We can define *practically adequate** beliefs as any believed proposition that is entailed by this strongest belief. So by brute force, practical adequacy* preserves single premise closure. Some may find this type of solution ad hoc because it's not clear why the strongest practically adequate belief has any normative significance.

Fortunately, we do have a principled explanation that comes from decision theory, but this requires us to think of categorical belief in a slightly different way. At the beginning of our discussion, I suggested that the practical role of belief was to take a proposition for granted in deliberation. I then explicated this idea within the decision-theoretic framework in terms of ignoring or eliminating columns in a decision table. On this view, the practical role of categorical belief is the following type of simplification. One at first considers an uncertain decision problem in all its complexity, but rather than reason on the basis of all these uncertainties, one simplifies by taking certain propositions as true or false.

There is, however, a slightly different simplifying role that categorical belief can play. On the constructive view of decision-making, beliefs may play a role in framing our decision problems, thereby helping us to demarcate a decision problem and construct the relevant deliberative judgments. If we take this view, when is a belief practically adequate? Here it is useful to borrow some terminology from L.J. Savage. Savage acknowledged that as creatures with bounded resources, we always engage in *small world* decision making. Here, a world is a particular representation of a decision problem rather than a possible world. While many factors could be relevant to a decision problem, we typically consider very little when thinking about and evaluating our choices. That is, we only consider a restricted set of acts, states, and consequences. Savage never offered a detailed picture of how we select small worlds but he did specify a condition on their selection. He argued that the preferences in a small world must be stable across refinements of that world to larger ones. In layman's terms, one has selected a good perspective on a decision problem only if what counts as rational from that perspective would remain stable when considering more and more possibilities and detail.

For example, suppose that you believe that it will not rain when deliberating about what to wear for the day. Framing your decision problem on this assumption, you conclude that it is best to wear a light

⁵⁷If there isn't a single strongest belief, then I'll assume that there is some way to select between the set of strongest beliefs.

sweater. This is a small world, a world in which one restricts what one considers. For Savage, this world is appropriate for use only if the same choice is best had one considered the likelihood of rain. So, if wearing a light sweater would still count as best from the perspective of the decision problem that accounts for these possibilities, then the original framing could be appropriate. Furthermore, Savage thought that a small world was appropriate only if what counts as rational remains stable across every possible refinement.

On this view, categorical belief does not simplify a decision problem that one has already considered. Rather, beliefs are used to frame a simplified decision problem. Now, appealing to this practical role of belief, we can define a new notion of practical adequacy. Consider the set of all propositions that are relevant to one's decision problem, and define *the smallest world* as the strongest proposition in that set that one may believe when framing one's decision problem such that the set of rational choices would remain stable under every relevant refinement.⁵⁸ A belief that P is *practically adequate*** just in case P is entailed by the smallest world. Smallest worlds are the most economical perspectives that one can take on decision problems without sacrificing anything of rational importance.

The two proposals may seem very similar, and it may simply be a matter of aesthetic judgment as to which is more natural. When stability across refinements is the standard for practical adequacy, I do find it natural to prioritize the most coarse-grained decision problem that possesses this property. After all, refinement stability is a counterfactual property and smallest worlds represent an optimal solution. However, when stability across column elimination is the standard for practical adequacy, it's less clear why we should prioritize the largest set of columns we can eliminate. Why isn't every appropriate elimination just as good? Whatever one's preferences between these two approaches, the important point is that the constructive decision-theoretic framework allows us to explore different roles for categorical beliefs and along with these different roles, different notions of practical adequacy.

4.2 Minimal Condition. As I noted above, the minimal condition that I have considered is very weak. And it may be odd that there are contexts in which we know, for example, that a fair coin will land heads on one of two tosses. Pragmatic epistemologies can certainly allow knowledge to have such features. However, it is also important to note that the proposed account can easily be modified to avoid such

⁵⁸Refinements are relevant just in case we are refining our decision problem with a possibility that is relevant to one's decision problem. If there isn't a unique smallest world, I'll assume that there is some way to select between the set of smallest worlds.

consequences. I started with the weakest minimal epistemic condition because that was all that was required to deal with Gettier cases. So I hoped to obey the Peircian creed to not place roadblocks to inquiry. More importantly, by starting with this condition, we gained some interesting insight into the nature of Gettier cases.

Nevertheless, I feel the pull of stronger conditions on knowledge. We could propose that one possess conclusive evidence for P or possess more than mere statistical evidence for P . One can even adopt a pragmatic infallibilism whereby, given what is deemed relevant in a context, one ought to be certain, remove all relevant doubts, or eliminate all alternatives. Thus, the specification of an appropriate minimal condition is another place for further exploration.

My aim here has not been to provide an unobjectionable pragmatic account of knowledge. Rather, I had two goals in mind. First, I wanted to lay out some general challenges for pragmatic epistemologies. Second, I hope to have offered a framework that is rich enough not only to address these particular challenges but many others as well. Thus, I hope to have showed that pragmatists need richer deliberative accounts. And once we appeal to these richer accounts, we will find that there is much more to explore within the pragmatic view of knowledge.

A Objective Relevance

To offer a more detailed account of objective relevance, let me introduce a distinction and some terminology. From the constructive point of view, there is a difference between bringing a possibility to mind and considering that possibility in one's deliberation.⁵⁹ I can bring to mind the possibility that I am a brain-in-a-vat but this does not mean that I have thereby taken this possibility into account when I deliberate. In order to take this skeptical possibility into account, I must reconstruct my beliefs over all the relevant states in a such a way that accounts for this possibility.⁶⁰ Just thinking about a possibility does not mean that one has accounted for it.

The notion of objective relevance is meant to capture the idea that there are some actualized states whose consideration would affect the probability of the previously considered states. To capture the idea

⁵⁹This distinction is different from the distinction between updating and imagining. The latter two only apply when the relevant proposition has already been considered.

⁶⁰This is not an update and is not equivalent to any type of conditionalization.

that probabilities are constructed relative to a space of relevant states, let $p^X(s)$ be the probability of s where $s \in X$ and X is the set of relevant states. Next, we must capture what it would be to take a previously unconsidered state z ($z \notin X$) into account. We can capture this by introducing the concept of refinement. Let me do so informally by way of example. Consider the coarse-grained set of states $X = \{\text{Rain, No Rain}\}$. We can refine X by taking the possibility of wind into consideration. The result will be the set of states $Y = \{\text{Rain and Wind, Rain and No Wind, No Rain and Wind, No Rain and No Wind}\}$. So if Y is a refinement of X that takes into account the unconsidered state y , $p^Y(s)$ is the newly constructed probability that has considered both y and $\neg y$. Now we can offer the following definition.

Objective Relevance: A state y is objectively relevant if and only if y is actual, $y \notin X$, and for some $s \in X$ where X is the set of subjectively relevant states and Y is the refinement of X that takes y into consideration, $p^Y(s|y) \neq p^X(s)$.

As I noted above, the probability can either be seen as a counterfactual subjective probability or as an evidential probability. Objectively relevant states are actual states of the world whose consideration as being true would affect probability (in one of these sense) over the set of previously considered states. For Catherine's decision problem, the state of the world in which there is an accident is objectively relevant. There are two competing explanations. It is objectively relevant because there has been an accident and if Catherine considered this possibility and took it for granted that y obtained, the resulting subjective probabilities would be affected. Or it is objectively relevant because there has been an accident and our assessment of Catherine's strength of evidence is affected by taking seriously the possibility that there may or may not have been an accident.

It is important to note that $p^Y(s|y) \neq p^X(s)$ does not entail that $p^Y(s|\neg y) \neq p^X(s)$.⁶¹ If it did, then almost every skeptical alternative and its negation might count as objectively relevant leaving us with an unrealistic account of knowledge. A simple example will explain why. Consider a state of the world in which Magneto, the villain of X-Men fame, has held up traffic all over the city. Assuming that such a state obtains, the probability that there is traffic should increase. However, if one takes this possibility seriously but assumes that it does not obtain, one should not be more or less confident that there is traffic

⁶¹This means that the types of conditional probabilities that I am considering are different from standard conditional probabilities. After all, typically $p(x|y) \neq p(x)$ if and only if $p(x|\neg y) \neq p(x)$. The reason the two types of conditional probabilities diverge is simple. Conditional probabilities assume that probabilities are defined over what one is conditionalizing upon. However, in the cases I am considering, we are conditionalizing on states that have not yet been considered so no probabilities have been defined (i.e., constructed).

on the freeway. When it comes to outlandish possibilities, it does not matter whether they are disregarded or taken for granted that they do not obtain. The probabilities remain the same either way.⁶²

One may try to reintroduce skeptical possibilities by conjoining them with true lottery-type possibilities.⁶³ For example, let y = Magneto has held up traffic or that Bill has not won the lottery. Is this possibility evidentially-relevant? Is $p^Y(\text{Traffic}|y) \neq p^X(\text{Traffic})$? I can only offer a rough response. My intuitive response is that practically all of the possibilities where y is true are those possibilities where Magneto has not held up traffic and Bill has not won the lottery. And the likelihood of traffic in those possibilities is determined by the likelihood of traffic when those possibilities are ignored.⁶⁴ So these gerrymandered possibilities are evidentially-irrelevant.

In addition, one might think that there are certain types of skeptical possibilities that are in fact true. For example, what about the possibility that someone is dreaming right now rather than the more direct skeptical possibility that I am dreaming. To my mind, possibilities like the former are not genuine skeptical possibilities since they are not perfectly compatible with my evidence. For that reason, it's unclear whether such possibilities are even evidentially-relevant. The generic dreaming hypothesis only counts as a genuine skeptical possibility if it were more specific – if someone were dreaming in a way that perfectly matches my current experiences. However, I take it that such a possibility is not actual.

Two additional objections present themselves. First, many objective facts are misleading and in accounting for them, we would have an inaccurate assessment of the DM's epistemic state. For example, if Catherine is unaware that the President is in town, that possibility is nevertheless objectively relevant for Catherine's assessment of the likelihood of traffic. However, it may also turn out that the President's staff has decided not to travel through the city in a motorcade, choosing instead to travel by helicopter. In this situation, the initial fact about the President whereabouts is misleading. And were Catherine to account for only this possibility, her belief that there is no traffic may be weakened enough to undermine

⁶²Taking the subjective interpretation of probability, there is a simple argument for why the shift from disregarding P or not P to believing that $\neg P$ should not always affect one's beliefs. There are a large number of outlandish skeptical possibilities. And if we raised our probabilities for every skeptical possibility (for which we possess a non-zero confidence) that was assumed not to hold, this would give us an artificial and unreasonable way for us to become extremely confident in contingent propositions.

⁶³Other gerrymandered possibilities are ruled out for alternative reasons. For example, true lottery propositions conjoined with the target proposition are inappropriate since they could not be used as genuine refinements of the original space of possibilities.

⁶⁴Increasing the likelihood of Bill's winning the lottery does not change matters since practically none of Bill's winning scenarios are scenarios where Magneto holds up traffic.

her claim to knowledge. In response, it is important to note that such facts are misleading only when considered on their own. And if all the objectively relevant facts are considered, we will be left with an objectively accurate assessment of the DM's epistemic state.

Finally, it may be objected that there are too many objectively relevant facts, resulting in a psychologically unrealistic account of knowledge. This objection does have some intuitive pull. In fact, for subjects who are experts about traffic, a very large number of states may be objectively relevant. After all, epistemically well-positioned subjects should be able to identify evidential connections that others would not be able to recognize. For example, experts may possess statistical information about the relationship between all types of weather and traffic making a multitude of facts about the weather relevant. Though I acknowledge the worry, I am not sure how problematic it is and can only offer a tentative reply. It should be admitted that in most situations, there are a great deal of objectively relevant states whether we interpret probabilities subjectively or evidentially. However, why should we think that there are an unmanageable number of them? Our boundedness as cognitive agents place an upper-bound on the number of evidential connections that we recognize. And if there is an upper-bound on what counts as evidence, there is an upper-bound on the number of objectively relevant states. Moreover, it seems plausible that we can and do make intuitive judgments about how strong a subject's epistemic state is given what is in fact true. It is exactly this type of objective epistemic evaluation that I am trying to capture.

References

- Anderson, Charity and John Hawthorne (forthcoming). "Knowledge, Practical Adequacy, and Stakes." In: *Oxford Studies in Epistemology*. OUP.
- Brown, Jessica (2008). "Knowledge and practical reason." *Philosophy Compass* 3.6, pp. 1135–1152.
- Chan, Timothy Hoo Wai (2014). *The aim of belief*. Oxford University Press.
- Coffman, E.J. (forthcoming). "Three approaches to Gettier Belief." In: *The Gettier Problem*. Ed. by Peter Klein Rodrigo Broges Claudio de Almeida. Oxford Univ Press.
- Douven, I. (2005). "A contextualist solution to the Gettier problem." *Grazer Philosophische Studien* 69.1, pp. 207–228.
- Drèze, J.H. (1990). *Essays on economic decisions under uncertainty*. Cambridge Univ Press.
- Fantl, J. and M. McGrath (2002). "Evidence, pragmatics, and justification." *The Philosophical Review* 111.1, pp. 67–94.
- (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Finlay, Stephen and Mark Schroeder (2015). "Reasons for Action: Internal vs. External." In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2015.
- Ganson, Dorit (2008). "Evidentialism and pragmatic constraints on outright belief." *Philosophical Studies* 139.3, pp. 441–458.

- Gilboa, Itzhak (2009). *Theory of Decision under Uncertainty*. Cambridge Books. Cambridge University Press.
- Greco, John (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Hawthorne, J. (2006). *Knowledge and lotteries*. Oxford Univ Pr.
- Hawthorne, J. and J. Stanley (2008). “Knowledge and action.” *Journal of Philosophy* 105.10, pp. 571–590.
- Hawthorne, John and Maria Lasonen-Aarnio (2009). “Knowledge and objective chance.” *Williamson on knowledge*, pp. 92–108.
- Heller, M. (1999). “The proper role for contextualism in an anti-luck epistemology.” *Nous* 33, pp. 115–129.
- Jeffrey, Richard (1965). *The Logic of Decision*. Chicago: University of Chicago Press.
- (1983). *The Logic of Decision, Second Edition*. Chicago: Univ of Chicago Pr.
- Kahneman, Daniel and Amos Tversky (1984). “Choices, values, and frames.” *American psychologist* 39.4, p. 341.
- Kim, Brian (2014). “The Locality and Globality of Instrumental Rationality: The Normative Significance of Preference Reversals.” *Synthese* 191.18, pp. 4353–4376.
- (2017). “Pragmatic encroachment in epistemology.” *Philosophy Compass* 12.5.
- Kripke, Saul A. (2011). “Two Paradoxes of Knowledge.” In: *Philosophical Troubles. Collected Papers Vol I*. Ed. by Saul A. Kripke. Oxford University Press.
- Lehrer, Keith and Thomas Paxson (1969). “Knowledge: undefeated justified true belief.” *The Journal of Philosophy* 66.8, pp. 225–237.
- Lewis, D. (1996). “Elusive knowledge.” *Australasian Journal of Philosophy* 74, pp. 549–567.
- Lichtenstein, Sarah and Paul Slovic (2006). *The Construction of Preference*. Cambridge University Press.
- Locke, Dustin (2015). “Practical Certainty.” *Philosophy and Phenomenological Research* 90.1, pp. 72–95.
- Plantinga, Alvin (1993). *Warrant and proper function*. Oxford University Press.
- Pritchard, D. (2007a). “Anti-luck epistemology.” *Synthese* 158.3, pp. 277–297.
- (2007b). *Epistemic Luck*. Oxford University Press.
- Raiffa, H. and R. Keeney (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley.
- Ross, Jacob and Mark Schroeder (2012). “Belief, credence, and pragmatic encroachment.” *Philosophy and Phenomenological Research*.
- Savage, L.J. (1972). *Foundations of Statistics*. New York: Dover.
- Shafer, G. (1986). “The construction of probability arguments.” *Boston University Law Review* 66.3-4, pp. 799–816.
- Shope, R.K. (1983). *The Analysis of Knowing: A Decade of Research*. Princeton University Press.
- Sosa, Ernest (2010). *Knowing full well*. Princeton University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford University Press, USA.
- Stroud, Barry (2000). *Understanding human knowledge: philosophical essays*. Oxford University Press, New York.
- Thomason, Richmond (1986). “The Context-Sensitivity of Belief and Desire.” In: *Reasoning about actions and plans*. Ed. by Michael Georgeff and Amy Lanksy. Morgan Kaufmann Publishers Inc., pp. 341–360.
- (2007). “Three Interactions Between Context and Epistemic Locutions.” In: *Modeling and Using Context*. Ed. by T. R. Roth-Berghofer D. C. Richardson B. Kokinov. Springer, pp. 467–481.
- (2014). “Belief, Intention, and Practicality: Loosening Up Agents and Their Propositional Attitudes.” In: *Epistemology, Context, and Formalism*. Springer, pp. 169–186.
- Tversky, Amos and Daniel Kahneman (1981). “The framing of decisions and the psychology of choice.” *Science* 211.4481, pp. 453–458.
- (1986). “Rational choice and the framing of decisions.” *Journal of business*, S251–S278.

Weatherson, Brian (2012). "Knowledge, bets, and interests." In: *New Essays on Knowledge Ascriptions*. Ed. by Jessica Brown and Mikkell Gerken. Oxford University Press.

Zweber, Adam (forthcoming). "Fallibilism, Closure, and Pragmatic Encroachment." *Philosophical Studies*, pp. 1–13.