

DECISION THEORY, PRAGMATIC ENCROACHMENT, AND GETTIER CASES

Brian Kim

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. Furthermore, I show that Gettier cases pose two problems that are specific to pragmatic accounts of knowledge. By addressing these problems through the development of a decision-theoretic account of knowledge, we can gain insight into what pragmatists think knowledge is. The result is a unique *lack of deliberative skill* explanation of why we lack knowledge in Gettier cases.

Central to epistemological orthodoxy is the belief 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 the role that knowledge plays in deliberation. And 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 principle along with a carefully crafted pair of cases offers an intuitive argument for pragmatic encroachment. For example, consider Catherine who possesses very strong but fallible evidence for 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 given Catherine's strong epistemic state, is she always justified in taking it for granted that she was born in NY? 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 it for granted that she was born in New York. 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 particular way, by determining the epistemic standards that must be met in order to know. So the reason

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 (c.f. Weatherson 2005).

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'.

why Catherine fails to know in the high-stakes scenario is because knowledge in this scenario requires that an extremely high standard be met.

My aim in this paper is not to evaluate the arguments for and against pragmatic encroachment but rather to explore its insights and consequences for our view of knowledge. If epistemological orthodoxy is to be abandoned and pragmatic encroachment embraced, what becomes of knowledge? What sort of relation is knowledge such that it can play its purported practical role? And how is the practical function of knowledge related to its other properties, such as its incompatibility with epistemic luck? In order to answer these questions, it will be useful to have a pragmatic analysis of knowledge. By doing so, we can gain insight into the type of relation that pragmatic encroachment theorists think knowledge is.

In recent years, proponents of pragmatic encroachment have tried to meet this challenge by proposing a number of different analyses. While a natural next step is to consider Gettier-style counterexamples, proponents of pragmatic encroachment have not discussed these cases in any detail.⁴ Given our aim of understanding the pragmatic view of knowledge, the Gettier problem presents some important challenges. In section 1, I show that Gettier cases bring to light some important problems that are specific to pragmatic accounts of knowledge. In section 2, I develop a pragmatic account that addresses these problems and show, in section 3, that this account can deal with Gettier cases. To conclude, I compare, in section 4, my preferred account with well-known alternatives. The upshot of our discussion is a pragmatic account where knowledge is the result of a type of deliberative skill. This account thereby offers a systematic explanation of when our beliefs fall short of knowledge. Unlike the ease of error or lack of credit explanations that are prominent in the contemporary literature, I propose a lack of deliberative skill explanation.

1 The Gettier Challenge for Pragmatists

In recent years, a number of pragmatic analyses 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 .”⁵ 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.”⁶ Weatherson proposes that “it is legitimate to write something on the decision table... iff the decision maker knows it to be true.”⁷

Unfortunately, each of these proposals fails to offer a satisfactory analysis. Fantl and McGrath’s proposal is unsuitable for a simple reason. 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 analysis does connect the concepts of knowledge and reason, we are nevertheless left without a clear account of the type of reason knowledge is.⁸

Weatherson’s proposal as well as Hawthorne and Stanley’s are unsuitable for one of two reasons. Either each analysis falls prey to Gettier-style counterexamples or, like Fantl and McGrath’s analysis, offers no explanation of Gettier intuitions, leaving us with unanswered questions about what knowledge is.⁹ Consider

⁴(Douven 2005) is an exception but he simply claims that Gettier cases are under-described, failing to identify the relevant practical factors.

⁵p.578, (Hawthorne and Stanley 2008)

⁶p.175, (Fantl and McGrath 2009)

⁷p.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.”

⁸These proposals flip traditional analyses on their heads by appealing to the reasons provided by knowledge rather than the reasons required for knowledge (e.g. Dretske 1971).

⁹(Brown 2008) uses Gettier situations as counterexamples to Hawthorne and Stanley’s analysis. (Neta 2009) and (Littlejohn

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. And, it 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.

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. The problem with this type of reply is that 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. We are then left with analyses that fail to provide insight into what knowledge is. Instead of analyzing knowledge in terms of reasons or in terms of rational decision-making, we are now gaining 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

Gettier cases also raise some specific problems for pragmatic accounts of knowledge. The subject's practical situation is irrelevant to our epistemic evaluations in Gettier cases. Since Catherine's belief that she will receive a small inheritance when she purchases her first house has been Gettiered, then she fails to know no matter what practical situation she is in. The irrelevance of practical factors poses a serious 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. However, in Gettier cases, we know that 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 the central pragmatic claim about knowledge.

The second problem arises from the fact that Gettier cases have taught us that knowledge is incompatible with epistemic luck.¹¹ 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 appear motivated in part by the need to provide a pragmatic and externalist account of knowledge. After all, on some views, what counts as a reason is independent of the internal state of the subject.¹³ 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 reverse the direction of analysis (using our intuitions about knowledge to analyze the nature of reasons). Unless we have an independent account of the type of reason that knowledge is or provides, the appeal to reasons will be unhelpful.¹⁴

2009) offer different criticisms of the proposed connection between knowledge and reasons.

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

¹¹While there are various types of epistemic luck (see Pritchard 2007b), the relevant type is veritic luck, which is present when it is lucky that a belief is true.

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

¹³(Finlay and Schroeder 2015)

¹⁴A detailed pragmatic account of knowledge appears to require a detailed account of deliberation. While I am not in

The decision-theoretic approach offers an alternative but faces serious problems as well. 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. Thus, for pragmatists such as myself who are partial to the decision-theoretic apparatus, the challenge is to show how we can develop an anti-luck epistemology within this framework.

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

1. We must show how to consistently claim that the practical situation of the subject determines the standards required for knowing 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.

On the account I will develop, the practical situation of the subject does in fact determine the epistemic standards that must be met for 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 develop an externalist version of decision theory. Therefore, I will address the first challenge by first addressing the second challenge. Before we turn to develop the pragmatic account, it is worth reiterating that I shall not be giving arguments for the proposed account. Rather, my aim is to simply develop an account that is able to address the challenges presented by Gettier cases.

2 A Decision-Theoretic Account of Knowledge

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 in which the notion of warrant is used to serve as a placeholder for the new condition, which is sufficient to turn true belief into knowledge.¹⁵ 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.

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.¹⁶ When we engage in belief-based reasoning, we engage

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.

¹⁵Warrant is the "elusive quality or quantity enough of which, together with true belief, is sufficient for knowledge." (v, Plantinga 1993)

¹⁶(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 premises in reasoning, acting on p, and simplifying one's decision by assuming that p. There are, however, some alternatives that are not equivalent. There is a difference between simplifying one's decision-making by assuming that p and p being assumed when framing one's decisions. One key difference is that our actual preferences are constructed relative to a framing of a decision problem. So the norms that govern the simplification of a decision problem can be related to constructed preferences. However, the norms that govern how to frame a decision problem cannot.

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.¹⁷ 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. One intuitive proposal is that such a simplification is warranted only if it makes no practical difference to one’s deliberations.¹⁸ 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. So let’s turn to consider the Bayesian view of rational deliberation, which my pragmatic analysis will be built upon.

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, one rationally deliberates just in case one coherently evaluates the choices that are available in a given decision problem.

To introduce the Bayesian view of a decision problem, consider the case where Catherine is deciding what mode of transportation to take on the way home from the museum. The Bayesian assumes that the decision maker (henceforth 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 assume that she considers the 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 is determined by two independent factors, her desire for the consequences and their likelihood. And since the Bayesian assumes that the consequences are thought to be certain given the performance of an act in a world where one of the states is actualized, the likelihood of the states determines the likelihood of each consequence given the performance

¹⁷See (Thomason 1986), (Thomason 2007), and (Thomason 2014) for further discussion.

¹⁸This 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.

of an act.¹⁹ Thus, Catherine’s evaluation of the acts depends only on her desire for the consequences and her belief about traffic.²⁰

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 constraints on deliberative preferences. One example of such a constraint is that rational preferences are transitive. If a DM prefers act A to B 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 a decision problem. 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.²¹ The probability and utility models function as a regulative ideal for our deliberative judgments since they represent all and only those 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.²²

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. 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. 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. 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 practical 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

¹⁹Some decision-theories are less-restricted by accounting for the conditional probability of the state given the act, thus violating act-state independence. However, as (Jeffrey 1976) shows, act-state independence can always be preserved.

²⁰Following Savage, I have separated the objects of belief and desire, which means that belief and desire are separable.

²¹More specifically, one’s beliefs, desires, and preferences are coherent if and only if they are respectively representable by a unique probability function, a utility function unique up to a positive linear transformation, and a probability weighted utility function.

²²(Savage 1972) offers the paradigmatic proof.

preferences.²³

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. 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 $EU(Train|\neg T) < EU(Cab|\neg T)$.²⁴ Her belief that $\neg T$ would also be unstable if we kept the utilities fixed and $p(T) > .4$.²⁵

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 always prefer the same choice.²⁶ 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 believe either P or $\neg P$ 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.²⁷ 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 judge that P is more likely than $\neg P$, meaning that $p(P) > .5$.²⁸ So S’s belief that P is pragmatically warranted only if $p(P) > p(\neg P)$ and S’s preferences given S’s belief that P are practically coherent with S’s ideal preferences.²⁹

²³Sets 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. Another criticism of pragmatic accounts is that they reject single premise closure for knowledge (see Zwebner forthcoming). So you may know a conjunction but not one of the conjuncts. While I shall not address this issue, I believe this issue boils down to the role that we think belief plays in decision-making. So the distinctions made in footnote 16 are relevant to this criticism.

²⁴ $EU(Cab) < .25(.2) + .75(.6) < .5$

²⁵ $EU(Cab) < .4(.2) + .6(.7) < .5$

²⁶For this reason, one might conclude that, in these decision problems, all-or-nothing beliefs have no role to play and so it is strange to ask what the subject knows. While I agree that there is something strange about knowledge of states in cases of dominance, I simply use the example to raise the general point that there is a minimum standard required for knowledge.

²⁷There 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.

²⁸Because the condition is stated in terms of probabilities, it requires that the subject’s degrees of belief are rational.

²⁹Though (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 $p(P) > p(\neg P)$ as a necessary for 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, while it seems clear that possessing some epistemic reason in

This condition will serve as the minimum epistemic standard for knowledge and failing to meet this standard would mean that one's belief is epistemically arbitrary. If one's belief is epistemically arbitrary, then one fails to know in every practical situation. In section 3, I will show that a particular type of epistemic arbitrariness can be used to account for our intuitions in Gettier cases. However, the condition as it has been described so far will not do. After all, subjects are reasonably confident that their Gettiered beliefs are true and so they satisfy this minimal condition. What we need and what I will develop is a decision-theoretic framework in which this minimal condition on probability is interpreted as a more objective measure of rational confidence, one that measures the strength of the subject's evidence given the actual situation they are in. To do so, we will need to develop a decision-theoretic framework where the probabilities are not merely subjective probabilities but more akin to evidential probabilities.

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 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 is a rich debate regarding coherence constraints on beliefs, desires, and preferences.³⁰ 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.

2.3.1 Two Limitations of Bayesian Decision Theory. The first problem is that the Bayesian offers no detailed account of how one does or should frame one's decision problem.³¹ It is simply assumed that the DM can do this and does so in a reasonable way. 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,

favor of believing p is necessary for warranted belief, it also seems that even if one did possess such a reason, one should possess some practical reason as well. After all, why adopt the simplifying attitude of categorical belief if doing so offers no practical benefits. One upshot of adding this condition is that belief would never be warranted in cases of dominance.

³⁰See (Gilboa 2009) for a summary of the debate.

³¹There 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 their decision problem. 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)

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. Should Catherine account for the weather as she evaluates her choices? Should she consider 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 and why does the DM describe her problem in the way she does?

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.³² 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.”³³ 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.³⁴

In order to address these two limitations, we must add to the Bayesian theory some account of how a DM frames her decision problem and constructs her deliberative judgments. I will call the resulting theory, a **Constructive Decision Theory**.³⁵ 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*. For her decision problem, Catherine 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 goal is that they may change from one deliberative context to the next. In another context,

³²See the articles in section 2 of (Lichtenstein and Slovic 2006).

³³(Lichtenstein and Slovic 2006), p.12

³⁴In (Kim 2014), I offer a detailed discussion of the challenges raised by choice behavior research for normative decision theory.

³⁵Some 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).

Catherine may have the goal of getting home in the least stressful, most aesthetically pleasing way. This in turn will identify a different 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.³⁶ This is only a rough sketch of the account and when I address Gettier cases in section 3, I will offer a more detailed discussion of the construction of probability.

The type of constructive decision theory that I have so far described remains Bayesian in spirit. Rational deliberative beliefs, desires, and preferences are still probabilities, utilities, and probability weighted utilities. The only change is that we incorporate how a DM sets 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 probability and utility judgments are based upon constructive judgments which 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.

³⁶This merely induces a partial ordering but if we also assume that goals also determine the relative importance of values then we may be left with a utility unique up to positive linear transformation.

So the demand that deliberative beliefs, desires, and preferences be representable by a unique probability, utility, and probability weighted utility is restricted to the decision problem for which these judgments are constructed.³⁷

2.3.3 Beyond coherence. While the constructive decision theory described in the previous section is still just a coherentist account of rational deliberation, the recognition that every comprehensive decision theory must account for how to set the parameters of a decision problem (i.e., 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 or objective factors can be appealed to in setting the parameters of a decision problem. By appealing to these objectively-framed decision problems, we will be able to develop an externalist epistemology.

In order to deal with Gettier cases, we will only need to reconsider how to determine which states are evidentially-relevant. And 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 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 own judgments. For example, consider the case in which Catherine takes into account the possibility that the President is visiting the city and 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.

In contrast, consider a different 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. However, let's suppose that Catherine judges that car accidents typically result in traffic. So Catherine herself judges that if there is an accident, then it is more likely that there is traffic. Since there has in fact been an accident, then from the point of view of the objective observer, she ought to consider this possibility.³⁸ Moreover, the explanation of its relevance appeals to both an objective fact as well as the DM's own probability judgments (i.e. subjective conditional probabilities). First, there has been a car accident. Second, from the point of view of the omniscient observer, given the information Catherine has, 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 subjective likelihood of any relevant state ought to be taken into account. While this intuitive presentation of objectively-relevant states will be sufficient for our discussion, I have provided a more detailed account in Appendix A.³⁹

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

³⁷In (Kim 2014), I explore how this very limited interpretation of BDT could be used to account for the more global aspects of instrumental rationality.

³⁸The set of states is a boolean algebra so if Catherine considers this possibility, she must also consider the possibility that there is no accident on the streets.

³⁹So far, I have only discussed how certain facts, unbeknownst to the subject, might be appealed to as an objective decision parameter that is used to demarcate the set of evidentially-relevant states. While it will go beyond the scope of this paper to discuss other objective decision 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.

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. We can now talk about rational degrees of belief for objectively-defined decision problems. Such degrees of belief are not beliefs that one actually has but rather ones that one is rational to construct given one’s evidence and provided that one accounts for certain possibilities. I shall offer a more detailed discussion of what these “objectively” rational degrees of belief are in section 3, but for now we can appeal to them used to produce an account of pragmatically warranted belief.⁴⁰

Pragmatically Warranted Belief: S’s belief that P is pragmatically warranted if and only if, in an objectively-framed decision problem, $p(P) > p(\neg P)$ and S’s preferences given S’s belief that P are practically coherent with S’s ideal preferences.⁴¹

Thus, knowledge is pragmatically warranted true belief.⁴²

3 Gettier Cases

The pragmatic account imposes a minimum standard for knowledge whereby the subject’s rational degree of belief that P must be greater than her rational degree of belief that $\neg P$ within an objectively framed decision problem. I previously indicated 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, let’s consider a Gettier case, doing so will also help us to understand what “objectively” rational degrees of belief are.

Consider the situation that arises in the 1999 remake of *The Thomas Crown Affair*. In the movie, the protagonist secretly steals Monet’s *San Giorgio Maggiore at Dusk* from the Metropolitan Museum of Art. In a seemingly magnanimous gesture, he offers to replace the stolen painting with one by another French Impressionist, Camille Pissaro. At the climax of the movie, the police try to foil Thomas Crown’s plan to return the stolen painting, but he accomplishes the task by pulling the fire alarm at the museum. As the sprinklers activate, the surface of the donated painting dissolves, revealing it to be a perfect forgery, meticulously water painted over the previously stolen Monet.

Thomas’ elaborate ruse places Catherine, the insurance agent sent to help the police investigation, in a Gettier situation. As she looks at the donated painting, she believes that she is standing in front of a painting by a French Impressionist, her belief is true, and she is in a very strong epistemic state with respect to this

⁴⁰As I will discuss, objectively rational degrees of belief are best interpreted as measures of evidential support – relative to a certain way of fixing what counts as evidence. Thus, the probability judgments that coincide with these objectively rational degrees of belief are best interpreted as evidential probabilities. While decision theory is typically concerned with subjective probability, the move to a constructive framework allows us to deal with objective constructions of the decision problem, which then allows us to talk of evidential probabilities.

⁴¹For our purposes, the decision problem counts as objectively-framed if the set of evidentially-relevant states account for evidentially-relevant matters of fact. However, there may be reasons to incorporate other objective parameters into the analysis of pragmatically warranted belief. 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.

⁴²On 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 the problems with multi-premise closure that (Hawthorne and Stanley 2008) raise for decision-theoretic accounts and the problems with single-premise closure that (Zweber forthcoming) raise more generally for pragmatic views.

proposition. After all, her perceptual faculties are working properly, she is very familiar with Pissaro and his paintings, and she has no reason to think that Thomas Crown would present the museum with a forgery. Nevertheless, Catherine does not know that she is in front of a painting by a French Impressionist. Her justification comes from her relation to the painted surface rather than the genuine Monet lying underneath. Since the latter makes her belief true, Catherine is lucky to have a true belief and luck is antithetical to knowledge.

The pragmatic account states that whether or not one knows depends upon the subject's deliberative context. However, since we are told nothing about Catherine's deliberative context, let us consider every possible deliberative context Catherine could be in – placing no restrictions on her potential goals and the set of relevant consequences. Nevertheless, there are some states that we must assume are relevant. First, since we are asking whether or not Catherine knows that P : she is standing in front of a painting by a French Impressionist, we assume that P and $\neg P$ are relevant. Next, though there may be a variety of subjectively relevant possibilities, they may be ignored since my explanation of the Gettier cases will apply whatever the case. Finally, there are two objectively relevant states that Catherine must consider. The painting that Catherine is looking at is actually a fake [F] and if she knew this, she would find it extremely unlikely that she was standing in front of a French Impressionist's painting. So the truth or falsity of F is relevant for assessing the likelihood of P . In addition, there is a genuine Monet underneath the fake [M] which, by Catherine's own lights, entails that P so M and $\neg M$ are evidentially relevant as well. Since both possibilities are actual and evidentially relevant, they must be taken into account. So the set of possible decision problems Catherine could be in is restricted by the fact that she must consider the possible truth and falsity of P , F , and M .⁴³

Given the set of relevant states, how should Catherine construct her probabilities? How confident may she be? To answer this question, let me briefly summarize a picture, taken from (Shafer and Tversky 1985), of how we should and often do construct our probabilities.

The weighing of evidence may be viewed as a mental experiment in which the human mind is used to assess probability much as a pan balance is used to measure weight. As in the measurement of physical quantities, the design of the experiment affects the quality of the result.⁴⁴

They go on to describe a variety of mental experiments that we can use in order to weigh evidence and assign probabilities. Consider a very simple example. As I am sitting at the park, I see a man walking and I wonder, what is the probability that he will step onto the road with his right foot rather than his left? To answer this, I ask myself, what would cause the man to perform one or the other action? I imagine that the only relevant cause is a reason or desire to do so. Considering my evidence, I conclude that I know nothing about him that would suggest that he desires to use one foot or the other. Thus, using my evidence to develop a simple causal model that would predict his behavior, I judge that the man's use of his right or left foot is produced at random and assign both events a .5 probability.

There are a variety of mental experiments that one could use and as Shafer and Tversky point out, the appropriateness of an experiment depends upon the case and evidence at hand. Nevertheless, they share some general and intuitive features. First, since probability judgments weigh evidence, they must be supported by evidence. Next, one's evidence should be used to fit the case at hand to some canonical example involving probabilities. When we want to assign quantitatively unique probabilities, the frequency, propensity,

⁴³That she is in the modern art wing of the museum and that the placard says that the painting is by the Camille Pissaro are all objectively relevant as are other possibilities. However, since my discussion would be unaffected by accounting for these possibilities, I will, for the sake of brevity, ignore them.

⁴⁴(Shafer and Tversky 1985), p. 309

and betting interpretations of probability supply some canonical examples.⁴⁵ The frequency interpretation proposes that we “compare our evidence to the scale of chances by asking how often, in situations like the one at hand, the truth would turn out in various ways.”⁴⁶ The propensity interpretation proposes that we interpret “the evidence in terms of a causal model and then [ask] about the model’s propensity to produce various results.”⁴⁷ The betting interpretation proposes that we “[assess] our willingness to bet in light of the evidence.”⁴⁸

Given this picture of the construction of probability, how should Catherine construct her probabilities? In Gettier cases, we fix what evidence the subject possesses so let us first fix Catherine’s evidence to consist in the following:

1. Her extensive knowledge of museums and art, particularly her knowledge of the French Impressionists.
2. Her perceptual evidence that the painting in front of her looks like one painted by Pissaro.

In order to evaluate what rational degree of belief - $p(P)$ - Catherine may have, she might engage in the following mental experiment in which she compares her case to a case with frequencies. She can consider a test where she is placed in front of all the paintings in the same museum. She can then judge the percentage of cases in which her judgment that P would be correct. Based upon her evidence, she would likely judge that she would be very reliable.

This experiment cannot, however, be used to determine the probability judgment that is relevant for assessing whether or not Catherine knows that P . In order to do so, we must also take into account the objectively relevant possibilities that the perceived painting is or is not a forgery. Moreover, it is a crucial feature of Gettier cases that Catherine possess no evidence for or against F . Though one might think that Catherine does have some evidence, this is just because the forgery possibility has not been sufficiently described. For even with Catherine’s extensive knowledge of museums and art, her evidence says nothing about whether the painting is such a perfect forgery that museum experts would have no way of detecting it. Gettier possibilities must always be described so that they are perfectly compatible with the evidence that the subject is presumed to have. In this respect, Gettier possibilities are like radical skeptical hypotheses except that they are constructed relative to a fixed set of evidence rather than all possible evidence. Once a skeptical hypothesis is taken seriously, it’s hard to imagine what evidence could rule it out. Similarly, for whatever evidence a Gettiered subject is presumed to possess, once a Gettier possibility is deemed relevant, it’s hard to imagine what evidence could be presented in favor of its truth or falsity.

Once the forgery possibility is deemed relevant, Catherine would have to design a new experiment to weigh her evidence. Since she has no evidence for or against the forgery possibility, Catherine can imagine a different test where it is randomly selected whether the painting in front of her is a perfect forgery or real. $p(P)$ is then determined by the frequency with which she would correctly judge that P . On the basis of her evidence, she imagines that she would do no better than chance, so $p(P) = .5$.⁴⁹ This would entail that, from the objective point of view, Catherine’s belief that P is epistemically arbitrary. Therefore, she fails to know that P in every deliberative context she could possibly be in.

⁴⁵Other types of canonical examples are possible. (Vasudevan 2013) shows that in some cases unique *a priori* probabilities can be derived on the basis of strong symmetry judgments. These examples may provide another type of canonical example.

⁴⁶(Shafer and Tversky 1985), p.316

⁴⁷(Shafer and Tversky 1985), p.316

⁴⁸(Shafer and Tversky 1985), p.316

⁴⁹In essence, we are imagining that Catherine assess the following: $p(P|E) = p(P|E \wedge F)p(F) + p(P|E \wedge \neg F)p(\neg F)$. To assess $p(F)$, she must consider her evidence and since we stipulated that the subject has no evidence that directly supports F or $\neg F$ (i.e. $p(F|E) = p(\neg F|E)$), Catherine should judge that $p(F) = p(\neg F) = .5$. Moreover, in Gettier cases, the conditional probabilities are fairly straightforward, $p(P|E \wedge F) \approx 0$ and $p(P|E \wedge \neg F) \approx 1$. Therefore, Catherine calculates that $p(P|E) = .5$ and thereby judges that $p(P) = .5$.

To address potential objections, a couple of points should be made. First, one might propose a different set of mental experiments by which Catherine would come to assign a probability greater than .5. After all, Catherine has lots of additional information about master forgers, the art world, and museums. Using the additional evidence, Catherine may be rational in having a high degree of confidence that P . While it is true that she has additional evidence, it will not make a difference in this case. For the sake of simplicity, I had described the case somewhat sparsely. But once we consider all the details, we will see that none of Catherine’s evidence can make a difference. In the movie, Thomas Crown has taken all the necessary steps to trick Catherine and the museum. While Catherine knows that there is a single individual capable of making a sufficiently good forgery, he is in jail so that possibility is ruled out. However, unbeknownst to Catherine, this master forger has a daughter who is just as talented and being a friend of Thomas Crown’s, she has been commissioned to produce the forgery. Once we take all the objectively relevant states into account, Catherine’s knowledge about museums and the art world provides neither evidence for nor against P . Since Catherine cannot be sure that she has been tricked, from the objective point of view in which these more detailed alternatives must be taken into account, her rational degree of belief in P must be the same as $\neg P$.

The dialectic is simple. One may object to my assessment of Catherine’s epistemic state by appealing to some evidence that underwrites a higher than .5 degree of belief that P . But if we have identified a real Gettier case, the world must be such that certain facts undermine every piece of evidence that the subject can have for or against P . And in the end, once all the relevant facts are taken into account, the subject has no reason to believe P rather than $\neg P$.

The pragmatic account of knowledge asks us to consider the strength of Catherine’s evidence for her belief given the actual situation she is in. We do this by considering how confident she may be in objectively demarcated decision problems (i.e., when every evidentially relevant fact has been taken into account). If Catherine’s belief has been properly Gettiered, then were Catherine to account for everything of relevance from the objective point of view, she would have no epistemic reason to believe that P rather than $\neg P$. Translated into our decision-theoretic framework, this means that her “objectively” rational degree of belief in P is .5. Since her belief is epistemically arbitrary, she fails to know in every practical situation she could be in. This externally evaluated type of epistemic arbitrariness is an arbitrariness that governs the subject’s belief given the actual strength of her evidence. So it might be better described as *evidential arbitrariness*, which underwrites our Gettier intuitions and identifies the type of luck that is present in all Gettier cases. From this objective point of view, one’s belief is true by luck in Gettier cases in just the same way that one’s guess that an actually fair coin lands heads, regardless of what one thinks about it, turns out to be correct. From the objective point of view, it is a matter of pure luck that such guesses are true and, similarly, it is a matter of pure luck that such beliefs are true.⁵⁰

To conclude this discussion, we can use the common structure of Gettier cases to generalize the explanation I have offered. First, there is the *Target Proposition* T that the subject fails to know and her belief is based upon her *Evidence* E .⁵¹ Next, the *Gettier Possibility* G is one that obtains, is evidentially relevant and is described in such a way that the subject has no evidence for or against G (i.e. $p(G|E) = p(\neg G|E)$). Though she bases her belief that T on evidence that is unreliable in her actual situation, the *Lucky Proposition* L ensures that her belief is true. The pragmatic account’s resolution appeals to the fact that “objectively” rational degrees of belief must be constructed on the basis of the evidence we stipulate her to have. Given the

⁵⁰Luck comes in degrees. It is much luckier that an amateur golfer gets a hole in one than a professional golfer. In cases of pure luck, success is not at all attributable to one’s skill.

⁵¹(Shope 1983) focuses on the role that a false lemma plays in the subject’s justification for her belief that T . This false lemma can be given a place in the way I describe Gettier cases but we can also just talk of the evidence the subject has without appeal to the false lemma that the evidence supports.

subject's evidence and the objective relevance of the Gettier possibility, she must judge that $p(T) = p(\neg T)$. And since T is always chosen so that $p(T|G) \approx 0$ and $p(T|\neg G) \approx 1$, we can calculate:

$$p(T) = p(T|G)p(G) + p(T|\neg G)p(\neg G) \approx .5^{52}$$

From the objective point of view, the Gettiered subject's evidence entails that $p(T) = p(\neg T)$ and as a result, her belief that T fails to meet the minimum standard required for pragmatically warranted belief and for knowledge. Regardless of what deliberative context she is in, she does not know that T .

4 Post-Gettier Analyses of Knowledge

The pragmatic account that I have proposed has similarities to both the indefeasibility and relevant alternative analyses of knowledge.⁵³ Given these similarities, I want to show how the account differs from and improves upon these well-known accounts.

The relevant alternatives account of knowledge runs into difficulties demarcating the set of alternatives that must be ruled out in order to count as knowing. As is evident 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 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 decision maker's 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.

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.⁵⁴ 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.⁵⁵

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

⁵²Remember that $p(G|E) = p(\neg G|E)$ and $p(T|E \wedge G) = p(T|\neg E \wedge G)$ so $p(E)$ does not matter for $p(T)$. We also need to calculate $p(L)$ though it's not clear how this is to be done. Since $p(T) = .5$ and $p(G) = .5$, $p(L) = \frac{1}{4}p(L|T \wedge G) + \frac{1}{4}p(L|T \wedge \neg G) + \frac{1}{4}p(L|\neg T \wedge G) + \frac{1}{4}p(L|\neg T \wedge \neg G)$. Though it's clear that $p(L|T \wedge G)$ is close to 1 and $p(L|\neg T \wedge G)$ is close to 0, what should the other conditional probabilities be? If we assign each conditional probability .5, then $p(L) = p(\neg L) = .5$. Though this may seem a bit strange at first, it nicely entails that there is no deliberative context where the Gettiered subject may stably believe that L or $\neg L$. Thus, she neither knows that L nor knows that $\neg L$.

⁵³My explanation of our Gettier intuitions does not require the full resources of decision theory. After all, I have simply argued that a certain type of epistemic evaluation could be defined, on the basis of which we could judge a gettiered subject's belief as evidentially arbitrary. And this is all that is required to provide the account that I have proposed. However, I have taken the time to situate this account with a suitable decision-theoretic framework to address the challenges that were raised in section 1 and to show how such an account fits naturally within a pragmatic, decision-theoretic account of knowledge.

⁵⁴(Lehrer and Paxson 1969)

⁵⁵See the discussion of defeasibility analyses in (Shope 1983).

between subjectively and objectively relevant states, and 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 relevant possibilities and facts. And the thought is that once all the objectively relevant possibilities are incorporated, no particular fact can be misleading.

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 that P is one whose subject S could easily have *believed* falsehoods similar to P, in ways similar to how S actually believes P."⁵⁶ 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.⁵⁷ 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."⁵⁸ 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.⁵⁹

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 identify some relevant values and appropriately evaluate outcomes relative to these values. The resulting account proposes that Gettiered beliefs are those that have not been skillfully formed in a deliberative situation. In fact, as I noted, these Gettiered beliefs are true as a matter of pure luck and so the success of the belief is not at all attributable to the decision maker's skill. On the pragmatic account I have offered, knowers are individuals who are deliberatively skillful in forming their belief, and knowledge is the outcome of this skill.

A Appendix

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. 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. 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 our constructed probability judgments. To capture the idea that probabilities are

⁵⁶p.6, (Coffman forthcoming)

⁵⁷See (Pritchard 2007b) and (Hawthorne and Lasonen-Aarnio 2009) for some proposals.

⁵⁸p.11, (Coffman forthcoming)

⁵⁹(Greco 2010), (Sosa 2010)

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 for the DM 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, $p^Y(s|y) \neq p^X(s)$.

Objectively relevant states are actual states of the world whose consideration as being true would affect the subjective probability judgments one constructs over the set of relevant states. For Catherine's decision problem, the state of the world in which there is an accident 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.

It is important to note that $p^Y(s|y) \neq p^X(s)$ does not entail that $p^Y(s|\neg y) \neq p^X(s)$.⁶⁰ 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. Having learned that such a state obtained, one's confidence in traffic should increase. However, having considered this possibility and assumed that it does not obtain, one should not be more or less confident that there is traffic on the freeway. When it comes to outlandish possibilities, it does not matter whether one disregards them or takes it for granted that they do not obtain. One's beliefs remain the same either way.⁶¹

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 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.

Next, 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

⁶⁰This 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).

⁶¹There 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 an uncountable number of outlandish skeptical possibilities. And if we raised our probabilities for every skeptical possibility that was assumed not to hold, this would give us an artificial and unreasonable way for us to become extremely confident in contingent propositions.

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. 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.
- Dretske, Fred (1971). "Conclusive Reasons." *Australasian Journal of Philosophy* 49.1, pp. 1–22.
- 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.
- (1976). "Savage's Omelet." *PSA: Proceeding of the Biennial Meeting of the Philosophy of Science Association* 1976, pp. 361–371.
- Kim, Brian (2014). "The Locality and Globality of Instrumental Rationality: The Normative Significance of Preference Reversals." *Synthese* 191.18, pp. 4353–4376.
- 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.
- Littlejohn, Clayton (2009). “On Treating Something as a Reason for Action.” *Journal of Ethics and Social Philosophy* 4, pp. 1–1.
- Neta, Ram (2009). “Treating something as a reason for action.” *Noûs* 43.4, pp. 684–699.
- 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.
- 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.
- Shafer, G. and A. Tversky (1985). “Languages and designs for probability judgment.” *Cognitive Science* 9.3, pp. 309–339.
- 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.
- 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.
- Vasudevan, Anubav (2013). “On the A Priori and A Posteriori Assessment of Probabilities.” *Journal for Applied Logic*.
- Weatherston, Brian (2005). “Can we do without pragmatic encroachment?” *Philosophical Perspectives* 19.1, pp. 417–443.
- (2012). “Knowledge, bets, and interests.” In: *New Essays on Knowledge Ascriptions*. Ed. by Jessica Brown and Mikkel Gerken. Oxford University Press.
- Zweber, Adam (forthcoming). “Fallibilism, Closure, and Pragmatic Encroachment.” *Philosophical Studies*, pp. 1–13.