

Attitudes Toward Epistemic Risk and the Value of Experiments

Abstract. Several different Bayesian models of *epistemic utilities* (see, e.g., [37], [24], [40], [46]) have been used to explain why it is rational for scientists to perform experiments. In this paper, I argue that a model—suggested independently by Patrick Maher [40] and Graham Oddie [46]—that assigns epistemic utility to *degrees of belief* in hypotheses provides the most comprehensive explanation. This is because this *proper scoring rule* (PSR) model captures a wider range of scientifically acceptable attitudes toward epistemic risk than the other Bayesian models that have been proposed. I also argue, however, that even the PSR model places unreasonably tight restrictions on a scientist's attitude toward epistemic risk. As a result, such Bayesian models of epistemic utilities fail as normative accounts—not just as descriptive accounts (see, e.g., [31], [14])—of scientific inquiry.

Keywords: Scientific experiment, Epistemic risk, Epistemic utility, Bayesianism, Categorical belief, Degrees of belief, Proper scoring rule.

1. Introduction

Performing experiments is an activity of central importance in science. Thus, explaining why it is rational for scientists to perform experiments is a critical project in the philosophy of science. A major success that has been claimed for the Bayesian approach to scientific inquiry is that Bayesian models of epistemic utilities provide such an explanation.¹ In this paper, however, I argue that any such Bayesian model places unreasonably tight restrictions on a scientist's attitude toward epistemic risk. Thus, I conclude that Bayesian models of epistemic utilities fail as normative accounts of scientific inquiry.

The classic Bayesian explanation is in terms of *pragmatic* utilities (see [18]). However, scientists principally have *epistemic* goals. As a result, several different Bayesian models of *epistemic* utilities have been used to explain why it is rational for scientists to perform experiments.² Most of these models assign epistemic utility to believing a hypothesis, rejecting a hypothesis, and withholding judgement (see, e.g., [23, pp. 465–467], [37, pp. 55–56], [40, pp. 114–117]). In addition to such *categorical belief* (CB)

¹For other successes for the Bayesian approach to scientific inquiry, see [13], [25], and [4].

²This is part of a larger trend of applying a consequentialist framework to epistemology (see, e.g., [15, p. 87], [8, p. 224], [47]).

models of epistemic utilities, there are a few Bayesian models that assign epistemic utility to *degrees of belief* in a hypothesis. Paul Horwich [24, pp. 127–129] was the first to propose such a degree of belief (DB) model of epistemic utilities.³ Patrick Maher [40, p. 113] subsequently suggested a refinement to the Horwich model. He suggested that epistemic utility should be assigned to degrees of belief in such a way that “scientists maximize expected cognitive utility by holding to the probabilities they actually have.” In other words, epistemic utility functions should be *proper scoring rules* (see [52, p. 28]). Graham Oddie [46] independently proposed a very similar model. Maher, however, ultimately rejected the proper scoring rule (PSR) model of epistemic utilities in favor of a categorical belief model.

In this paper, I argue that the PSR model captures a wider range of scientifically acceptable attitudes toward epistemic risk than the other Bayesian models of epistemic utilities that have been proposed. These other Bayesian models are essentially restricted versions of the PSR model. I also respond to several important objections (principally those raised by Maher) to the PSR model.⁴ I conclude that the PSR model provides the most comprehensive Bayesian explanation for why it is rational for scientists to perform experiments.

But focusing on the PSR model also serves to highlight a serious flaw with all Bayesian models of epistemic utilities. The standard objection to the Bayesian approach to scientific inquiry is that it is not *descriptively* adequate. Real scientists are not Bayesian agents (see, e.g., [31, p. 450], [14, pp. 149–157]). In particular, scientists do not always have coherent degrees of belief and do not always update their degrees of belief by conditionalization. Even so, the Bayesian can always respond to this objection by claiming that scientists *ought* to be Bayesian agents (cf. [25, pp. 292–295]). But unfortunately for the Bayesian, the Bayesian models of epistemic utilities turn out to have extremely implausible *normative* implications. In this paper, I show that, in the special case where a scientist seeks the “truth and nothing but the truth,” there is only one attitude toward epistemic risk that is allowed by the PSR model. While this particular attitude is probably scientifically acceptable, there is no obvious reason to think that it is scientifically required. Since the PSR model is the least restrictive Bayesian model of epistemic

³Horwich [24, p. 127] actually talks about the “error in a probability valuation” rather than the epistemic utility of a degree of belief, but that is what his proposal amounts to (cf. [40, p. 112], [55, p. 52])

⁴The literature on Bayesian models of epistemic utilities is not very *well connected* (cf. [3, p. 11]). For example, Oddie does not refer to Maher or respond directly to his objections to the PSR model. See [47] for a fairly comprehensive review of the literature.

utilities, I conclude that these Bayesian models place unreasonably tight restrictions on a scientist's attitude toward epistemic risk.

2. The Goals of Scientific Inquiry and Preferences over Epistemic Lotteries

Whether it is rational for us to engage in a particular activity depends on whether doing so furthers the goals that we have. Thus, in order to explain why it is rational to perform experiments, we need to know what the goals of scientific inquiry are. A standard suggestion is that scientists want to discover *truths* about the world (cf. [15, p. 245], [37, p. 49], [46, p. 537]). In other words, scientists want to increase their degree of confidence in scientific hypotheses that are true.⁵ There are two main ways that this suggestion can be cashed out (cf. [15, p. 88]). In terms of *categorical belief*, scientists prefer to believe a true hypothesis than to withhold judgment on it or reject it (cf. [37, p. 55]). In terms of *degrees of belief*, scientists prefer to have a high degree of belief in a true hypothesis (cf. [15, p. 90]).

Of course, it is not always rational for scientists to perform experiments *all things considered*. There are often significant non-epistemic costs (e.g., ethical or financial ones) to performing an experiment that outweigh the epistemic benefits (cf. [41, p. 173]). But even if we restrict our focus to epistemic consequences, performing an experiment can have bad consequences. Performing experiments *typically* leads scientists to increase their degree of confidence in the truth. However, unless an experiment is 100 percent accurate, it is possible for a scientist to get misleading results. Thus, performing an experiment can sometimes lead a scientist to *decrease* her degree of confidence in the truth.

Performing an experiment is essentially like buying a lottery ticket that has epistemic prizes. Prior to performing the experiment, there is no guarantee that the epistemic benefits will outweigh the epistemic costs. Even so, we can still try to show that a scientist should always prefer this epistemic lottery to simply maintaining her current degree of confidence in the truth.⁶ As Maher [40, pp. 104–105] puts it, we want “some argument to show that

⁵Some philosophers (e.g., Sandra Harding [20, pp. 585–587]) have argued that increasing one's degree of confidence in the *truth* is not a goal of scientific inquiry at all. But see [15, pp. 3–40] for an extended defense of true belief as our epistemic goal.

⁶Epistemic lotteries are unlike most other lotteries in that one rarely (if ever) finds out for sure whether one has won the prize. Even so, it is still reasonable for a scientist to seek to increase her degree of confidence in the truth (i.e., to try to win the prize) and to have preferences over epistemic lotteries (cf. [55, p. 57], [32, p. 445]).

the chance of getting closer to the truth by gathering evidence outweighs the risk of getting further away from the truth.”

Whether or not it is worth the risk to buy a lottery ticket (epistemic or otherwise) depends on an individual’s attitude toward taking risks. And different individuals may have very different attitudes toward epistemic risk (cf. [37, p. 56], [36, p. 71], [10, p. 565], [35, p. 264], [55, p. 61], [8, p. 234]). David Hume [26, p. 111], for example, appears to have been extremely averse to epistemic risk while William James [28, pp. 31–32] appears to have been much more willing to take epistemic risks. Given that different individuals can have different attitudes toward epistemic risk, it is not immediately clear that it is always rational to perform an experiment. In fact, an individual might conceivably be so averse to epistemic risk that she prefers maintaining her current degree of confidence in the truth to performing any experiment that might lead to a decrease in her degree of confidence in the truth.⁷ However, we might try to argue that such extreme aversion to epistemic risk is *unscientific*.

An attitude toward epistemic risk is only *scientifically acceptable* if it is consistent with certain core scientific principles (cf. [41, pp. 209–216]). For example, a scientist must prefer true belief to error. Also, a scientist must only change her degree of belief in a hypothesis on the basis of new evidence (cf. section 4). Several Bayesian models of epistemic utilities have been used to show that it is always rational for an individual whose attitude toward epistemic risk is consistent with these principles to perform experiments. In the next section, I describe the basic framework for such Bayesian models.

3. Epistemic Utility Functions and the Value of Experiments

If an individual’s preferences over a set of lotteries meet certain “consistency” conditions (such as *transitivity* and *continuity*), then these preferences can be captured by a *utility function* (cf. [33, pp. 132–133], [6, pp. 504–505]). If an individual’s preferences over a set of *epistemic* lotteries (e.g., experiments) meet these same conditions, then these preferences can be captured by an *epistemic utility function* (EUF). For each doxastic state (e.g., believing a particular hypothesis) and each hypothesis, an EUF gives the epistemic utility of that doxastic state when that hypothesis is true. Let $H = \{h_1, \dots, h_n\}$

⁷This only occurs if any lower degree of confidence in the truth is worse than her current degree of confidence and no higher degree of confidence is any better than her current degree of confidence. However, even if her attitude toward epistemic risk is independent of her current degree of confidence in the truth, it is still possible to be so averse to epistemic risk that it is not rational to perform many experiments (see section 6).

be a set of mutually exclusive and jointly exhaustive hypotheses. $u_i(DS)$ will refer to the epistemic utility of doxastic state DS when h_i is true.⁸

An individual's utility function is defined in such a way that the expected utility of one lottery is greater than the expected utility of another lottery if and only if this individual prefers the first lottery to the second. In other words, a utility function captures this individual's attitude toward risk.⁹ As a result, we can just talk in terms of expected utility rather than talking directly about preferences over lotteries. In particular, it is rational for a scientist to perform an experiment if and only if the expected epistemic utility of performing the experiment is at least as great as the expected epistemic utility of not performing the experiment. In this paper, I will refer to the difference between these two values as the *epistemic value of performing the experiment*.¹⁰

The expected epistemic utility of not performing the experiment is simply the expected epistemic utility of a scientist's current doxastic state. Of course, since scientists almost never know *for sure* which hypothesis is true, doxastic states themselves are epistemic lotteries. Let DS be a scientist's current doxastic state and let $\mathbf{r} = (r_1, \dots, r_n)$ be her current probability assignment on H such that r_i is the probability she assigns to h_i being true. The expected epistemic utility of her current doxastic state is then $EU(DS, \mathbf{r}) = \sum_i r_i u_i(DS)$.

If the scientist performs the experiment, she could end up in one of several different doxastic states depending on what the result of the experiment is. The expected epistemic utility of performing the experiment is a weighted average of the expected epistemic utilities of these doxastic states. These expected epistemic utilities are weighted by how likely the scientist thinks that each of the possible results are. Let DS_e and \mathbf{r}_e be the scientist's

⁸Of course, any model of epistemic utilities will be an idealization (cf. footnote 24). In actual practice, scientists will rarely have an explicit list of mutually exclusive and jointly exhaustive hypotheses or clear preferences over all epistemic lotteries involving these hypotheses.

⁹Attitudes toward risk are sometimes discussed in terms of various decision rules that an individual might use. For example, an individual who is risk averse may tend to use the *maximin* rule rather than try to maximize her *expected* payoff. However, attitudes toward risk can also be discussed (as they will be in this paper) solely in terms of the utilities assigned to possible outcomes by an individual who always tries to maximize expected utility.

¹⁰This is also the standard technique for calculating the pragmatic value of performing an experiment (see, e.g., [18], [51, pp. 87–106]). In that case, we would simply be evaluating *acts* rather than *doxastic states*.

new doxastic state and new probability assignment if she gets result e .¹¹ $EU(DS_e, \mathbf{r}_e)$ will then be the expected epistemic utility of the scientist's new doxastic state if she gets result e . Finally, let $pr(e)$ be the probability that the scientist assigns—before performing the experiment—to getting result e .¹² The expected epistemic utility of performing the experiment is then $\sum_e pr(e)EU(DS_e, \mathbf{r}_e)$ where we sum over all possible experimental results.

Of course, in order to calculate the expected epistemic utility of performing the experiment, we still need to know one more thing. Namely, we need to know how to determine exactly what DS_e and \mathbf{r}_e are. The standard Bayesian assumption is that scientists get to \mathbf{r}_e from \mathbf{r} via *conditionalization* (cf. [18, p. 319], [24, p. 128], [40, p. 110], [46, p. 533]). But how we determine what DS_e is depends on exactly what counts as a doxastic state (e.g., on whether epistemic utility is assigned to categorical beliefs or to degrees of belief). As a result, I will defer this issue until I discuss the different Bayesian models of epistemic utilities.

Finally, as noted above, individuals may differ in their attitudes toward epistemic risk. Even among scientists, there may not be one single attitude toward epistemic risk (cf. [37, p. 64], [41, p. 210]).¹³ In other words, there may be more than one scientifically acceptable EUF. A complete explanation of why it is rational for scientists to perform experiments must establish that the epistemic value of performing *any* experiment is greater than or equal to zero for *all* scientifically acceptable EUFs. In the next section, I describe a *model of epistemic utilities* (i.e., a set of EUFs) that seems to provide such an explanation.

4. The Proper Scoring Rule Model

In any DB model of epistemic utilities, a scientist's probability assignment \mathbf{r} is her doxastic state. Thus, $u_i(\mathbf{r})$ is the epistemic utility of her doxastic state if h_i is true and $EU(\mathbf{r}, \mathbf{r})$ is the expected epistemic utility of her doxastic state. In addition, this means that DS_e simply is \mathbf{r}_e . Thus, conditionalization by itself allows us to determine what DS_e and \mathbf{r}_e are.

The PSR model, in particular, places only one constraint on EUFs: $EU(\mathbf{r}, \mathbf{r}) \geq EU(\mathbf{q}, \mathbf{r})$ for all \mathbf{q} . In other words, an EUF should be a *proper*

¹¹We also have to assume that she knows that DS_e and \mathbf{r}_e will be her new doxastic state and new probability assignment (cf. [51, p. 91]).

¹²This can be determined from \mathbf{r} together with the conditional probability that the scientist assigns to each possible experimental result given each hypothesis (cf. [8, p. 226]).

¹³For example, it is certainly the case that mathematicians are more averse to epistemic risk than other scientists (cf. [21, p. 295]).

scoring rule. This constraint only rules out EUFs that are clearly unscientific. If an individual's preferences over epistemic lotteries are not captured by a proper scoring rule, then pursuing her epistemic goals will sometimes require her to engage in the scientifically objectionable practice of jumping to conclusions.¹⁴ More specifically, she can increase her expected epistemic utility simply by changing her probability assignment from \mathbf{r} to some \mathbf{q} such that $EU(\mathbf{q}, \mathbf{r}) > EU(\mathbf{r}, \mathbf{r})$. As Maher [40, pp. 112–113] notes, “we think that such a shift, not motivated by any new evidence, does *not* in fact increase expected cognitive utility.”

While it may not be psychologically possible for scientists to simply change their probability assignment in the absence of new evidence, this is only a contingent matter (cf. [46, p. 536]).¹⁵ For example, we might develop a drug that makes us certain of our favorite scientific hypotheses. As a result, a scientifically acceptable EUF needs to be consistent with what we know about the goals of scientific inquiry. In particular, a scientifically acceptable EUF should be consistent with the principle—which Oddie [46, p. 538] refers to as *conservatism*—that “absent any new information you should not change your cognitive state.”

The PSR model has a number of virtues. First, it provides an explanation for why it is rational for scientists to perform experiments. Maher [40, p. 113] established that, if a scientist conditionalizes on the results of an experiment, then the epistemic value of performing an experiment is always greater than or equal to zero for all EUFs in the PSR model (cf. [41, p. 180], [46, pp. 539–541]).¹⁶

Second, the PSR model does not assume that it is rational for scientists to perform experiments. It only assumes that a scientist cannot increase her expected epistemic utility simply by changing her probability assignment. As Goosens [19, p. 96] and Maher [41, pp. 177–178] suggest, we do not want

¹⁴Analogously, we would not think that we had identified a *morally acceptable* goal if it were shown that pursuing this goal requires us to engage in morally objectionable practices. This is the form of many popular objections to utilitarianism (see, e.g., [43, pp. 468–469]).

¹⁵Several other epistemologists and philosophers of science (see, e.g., [41, pp. 147–148], [55, p. 57], [53, pp. 162–163], [8, p. 236]) have argued that models of epistemic utilities can be used to evaluate the rationality of our beliefs even if they are not under our direct voluntary control.

¹⁶Brian Skyrms [51, pp. 92–106] shows various ways in which this sort of result can be extended. For example, the epistemic value of performing an experiment is still greater than or equal to zero even if the scientist updates her probability assignment by *probability kinematics* rather than by strict conditionalization.

our explanation for why it is rational for scientists to perform experiments to beg the question.

Finally, with only this one constraint on EUFs, the PSR model captures many different attitudes toward epistemic risk (cf. [9, p. 179]). For example, two commonly discussed proper scoring rules are the *Brier rule*: $u_i(\mathbf{r}) = 2r_i - \sum_k r_k^2$ for all i (see Figure 1) and the *logarithmic rule*: $u_i(\mathbf{r}) = \log(r_i)$ for all i (see [52, p. 30]).¹⁷ In addition, as I argue in the next section, the PSR model captures a wider range of scientifically acceptable attitudes toward epistemic risk than any of the other Bayesian models of epistemic utilities. Thus, it provides the most comprehensive Bayesian explanation for why it is rational for scientists to perform experiments.

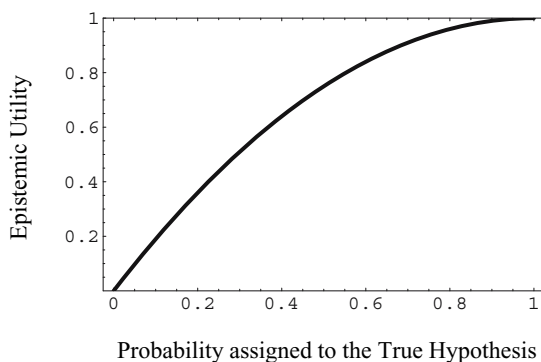


Figure 1. The Brier Rule

5. Other Bayesian Models of Epistemic Utilities

5.1. Categorical Belief Models

The first CB models of epistemic utilities were proposed by Carl Hempel [23, pp. 465–467] and Isaac Levi [37, pp. 55–56]. In these models, believing a true hypothesis is assigned the maximum epistemic utility, rejecting a true hypothesis (i.e., believing a false hypothesis) is assigned the minimum epistemic utility, and withholding judgment is assigned an epistemic utility somewhere between the two extremes.¹⁸

¹⁷Maher [40, p. 113] uses the Brier rule to show that the PSR constraint is satisfiable (cf. [46, p. 538], [8, p. 233], [34, pp. 385–387]).

¹⁸Strictly speaking, Hempel and Levi assign epistemic utility to *accepting* a hypothesis (cf. [40, pp. 114–117]). However, a number of philosophers of science (see, e.g., [29, p. 245], [46, p. 534], [53, pp. 161–162]) have argued that scientists do not accept hypotheses.

Maher [40, pp. 114–117] used a CB model to explain why it is rational for scientists to perform experiments.¹⁹ In this subsection, I show that CB models capture a wide range of scientifically acceptable attitudes toward epistemic risk. I argue, however, that the PSR model also captures these same attitudes toward epistemic risk. In addition, I argue that the PSR model captures scientifically acceptable attitudes toward epistemic risk that are not captured by CB models.

Levi's original model covers the simplest case where H consists of just two mutually exclusive and jointly exhaustive hypotheses (i.e., where h_2 is the negation of h_1).²⁰ The EUFs in this model have the following form: $u_i(B_i) = 1$, $u_i(W) = k$ such that $1 \geq k \geq 0$, and $u_i(B_{j \neq i}) = 0$ for all i . In addition, Levi assumes that a scientist's doxastic state is a function of her EUF and her probability assignment. In particular, he assumes that a scientist will always be in the doxastic state which has the highest expected epistemic utility. Together with conditionalization, this assumption allows us to determine what DS_e and \mathbf{r}_e are.

In Levi's original model, every possible value of k represents a different attitude toward epistemic risk (cf. [37, p. 56]). The higher the value of k , the more risk averse a scientist is. As a result, the model captures a wide range of preferences over epistemic lotteries. This is illustrated by the fact that the model captures the diverse attitudes of Hume and James. According to Hume [26, p. 111], "there is a degree of doubt, and caution, and modesty, which, in all kinds of scrutiny and decision, ought for ever to accompany a just reasoner." This suggests that it is typically better to withhold judgment than to believe a hypothesis that might turn out to be false. Hume's attitude toward epistemic risk can arguably be captured by setting k to a value very close to 1. According to James [28, pp. 31–32], "a rule of thinking which would absolutely prevent me from acknowledging certain kinds of truth if those kinds of truth were really there, would be an irrational rule."²¹ This suggests that it is always better to believe a hypothesis (even if that

In this paper, I will continue to talk about belief rather than acceptance since nothing here turns on this distinction.

¹⁹The CB model that Maher [40, p. 110] actually uses is somewhat complicated so that it can handle cases like "observations in quantum mechanics, since here there are no states which are both independent of the observation and also uniquely determine the outcome of the observation." I will simplify his model since nothing in this paper turns on this issue.

²⁰This model has been significantly revised and extended in his later work (see, e.g., [38], [39]).

²¹James specifically has in mind here truths that are "momentous." He [28, pp. 25–26] thinks that an aversion to epistemic risk can be appropriate in more mundane matters.

hypothesis might turn out to be false) than to withhold judgment. James' attitude toward epistemic risk can arguably be captured by setting k to a value close to 0.5.²²

There are also a couple of ways in which Levi's original model can easily be extended to capture an even wider range of preferences over epistemic lotteries. First, u_1 does not have to have the same shape as u_2 in a CB model.²³ Levi [37, p. 55] was explicitly trying to model the case where a scientist seeks the "truth and nothing but the truth." In that context, the epistemic utility of believing a true hypothesis should be the same regardless of which hypothesis happens to be true (cf. [41, p. 214]). However, scientists do not always seek the "truth and nothing but the truth." For example, $u_1(B_1)$ might be greater than $u_2(B_2)$ if h_1 is more informative, simpler, or has more explanatory power than h_2 (cf. [23, p. 467], [55, p. 55]). A CB model need only require that believing the true hypothesis is epistemically better than not believing it. In other words, EUFs must simply satisfy the following (*monotonicity*) constraint: $u_i(B_i) \geq u_i(W) \geq u_i(B_{j \neq i})$.

Second, there can be more than just two hypotheses in a CB model (cf. [38, p. 34]). If there are more than two hypotheses, there are just more possible doxastic states. For example, if there are *three* mutually exclusive and jointly exhaustive hypotheses, there are seven possible doxastic states. Namely, a scientist could believe one of the hypotheses and reject the other two (three possible states here), the scientist could withhold judgment on all three hypotheses, or the scientist could withhold judgment on two of the hypotheses and reject the third (three possible states here).²⁴

In addition, all of the preferences over epistemic lotteries captured by CB models seem to be scientifically acceptable. In other words, an individual who simply values believing true hypotheses (and not believing false hypotheses) does not appear to be unscientific. For example, all EUFs in CB models are consistent with conservatism (cf. [40, p. 116]).²⁵

²²If k less than 0.5, then it is always better for a scientist to believe some hypothesis than to withhold judgment (cf. [37, p. 55]).

²³For the same reasons, u_1 does not have to have the same shape as u_2 in a DB model either.

²⁴There are several other possible extensions of Levi's original model. For example, scientific hypotheses are not always mutually exclusive. Also, scientific hypotheses are not necessarily just true or false period (e.g., they can arguably have different degrees of *verisimilitude*). Finally, scientists are typically interested in having correct answers to many different questions. None of these issues affect any of the results discussed in this paper, however. Maher [41, pp. 141–142] shows how to handle these issues in a CB model and they can be handled in the same way in a DB model.

Philip Percival [47, p. 145] has suggested that there is a serious conflict between the CB models and the PSR model. However, while the PSR model certainly disagrees with the CB models about what counts as a doxastic state, the PSR model does not disagree with the CB models about what counts as a scientifically acceptable attitude toward epistemic risk. As I argue below, the PSR model captures all of the preferences over epistemic lotteries that are captured by CB models.

Any EUF in a CB model can be translated into a DB version.²⁶ First, we determine which doxastic state in the CB version maximizes expected epistemic utility when the scientist’s probability assignment is \mathbf{r} . Second, we take $u_i(\mathbf{r})$ to be the epistemic utility of that doxastic state when h_i is true. So, for example, in the case of Levi’s original model, $u_i(\mathbf{r}) = \{1 \text{ if } r_i \geq z; 0 \text{ if } r_i \leq 1 - z; z \text{ otherwise}\}$ for all i where $z = \max\{k, 0.5\}$ (see Figure 2 for the case where $k = 0.8$).

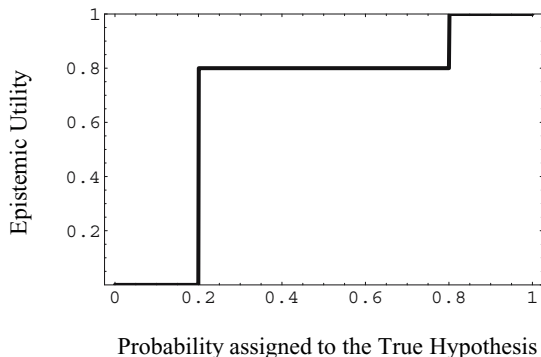


Figure 2. Levi’s step function

²⁵In a CB model, a scientist might decide to change her *doxastic state* in the absence of new evidence. For example, she might adopt a new EUF if she changes her mind about the importance of simplicity or about the simplicity of particular hypotheses. With this new EUF, a different doxastic state might have a higher expected epistemic utility. However, she will never be motivated to change her *probability assignment* in the absence of new evidence.

²⁶Goldman [15, p. 88] suggests that there might be such a translation between the CB and DB models, but does not provide one. This translation is reminiscent of the way in which James Hawthorne and Luc Bovens [22, p. 252] “model qualitative doxastic states in terms of a threshold value for a degree-of-confidence function.” See [12], [27], [32, pp. 447–457], [5, pp. 12–32], and [54] for other discussions of the connection between categorical belief and degrees of belief. Finally, it should be noted that there may be ways to assign epistemic utilities to categorical beliefs for which there is no such translation. However, as far as I know, no such Bayesian model of epistemic utilities has been put forward (or used to establish the epistemic value of performing experiments).

On this translation, the expected epistemic utility of a probability assignment \mathbf{r} in the DB version is the same as the expected epistemic utility of the scientist's doxastic state in the CB version when her probability assignment is \mathbf{r} . As a result, we get exactly the same result when we calculate the epistemic value of performing an experiment. Thus, the translated EUF preserves the scientist's preferences over experiments (i.e., epistemic lotteries). In addition, since a scientist cannot increase her expected epistemic utility simply by changing her probability assignment in a CB model, the translated EUF satisfies the PSR constraint.²⁷

It should be noted that I am not making any claim here about the metaphysical relationship between categorical beliefs and degrees of belief. In particular, I am not claiming that categorical beliefs are reducible to—or that they supervene on—degrees of belief (contrast [5, pp. 96–100]). I am only claiming that the PSR model captures all of the scientifically acceptable attitudes toward epistemic risk that CB models capture. For example, suppose that scientist **C** assigns epistemic utilities to categorical beliefs. Also, suppose that scientist **D** assigns epistemic utilities to degrees of belief, but that she does so according to the translation of **C**'s EUF. These two scientists may very well be assigning epistemic utilities to two metaphysically different sorts of things. Nevertheless, these two scientists have the very same preferences over epistemic lotteries. In other words, they have the same attitude toward epistemic risk.

In addition to capturing all of the preferences over epistemic lotteries that are captured by CB models, as I argue below, the PSR model also captures scientifically acceptable preferences over epistemic lotteries that are not captured by CB models. In particular, CB models do not capture the preferences of a scientist who thinks that she is in a better epistemic situation if she increases her degree of belief in a true hypothesis even if she does not come to believe that hypothesis or she already believed that hypothesis (cf. [15, p. 90]).

With any model of epistemic utilities, epistemic utility can only increase if a scientist's doxastic state changes. As a result, the epistemic value of performing an experiment will only be positive if it has the potential to change her doxastic state. Many experiments that can change a scientist's degree of belief will not be sufficiently decisive to change her doxastic state in a CB model (e.g., lead her to believe a hypothesis that she had been

²⁷For example, in Levi's original model, $EU(\mathbf{q}, \mathbf{r})$ is equal to either r_1 , r_2 , or z for all \mathbf{q} . $EU(\mathbf{r}, \mathbf{r})$, in particular, is equal to the maximum of these three values. Thus, $EU(\mathbf{r}, \mathbf{r}) \geq EU(\mathbf{q}, \mathbf{r})$ for all \mathbf{q} .

withholding judgment on). In a CB model, a scientist will be indifferent between performing such experiments and simply maintaining her current degree of confidence in the truth. For example, suppose that $k = 0.8$ in Levi’s original model and that a particular experiment \mathbf{X} is three times more likely to yield one result if h_1 is true and three times more likely to yield another result if h_2 is true. The epistemic value of performing \mathbf{X} is zero for a wide range of prior probabilities (see Figure 3).

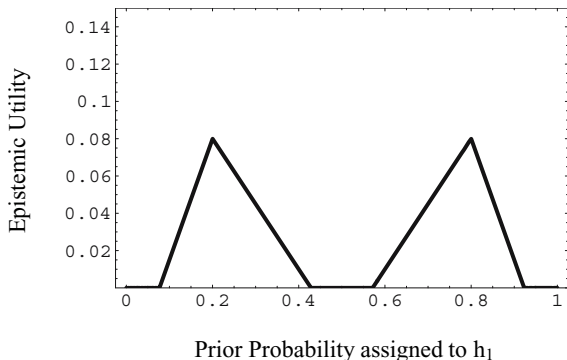


Figure 3. Value of an experiment, using Levi’s step function

However, a scientist who thinks that she is in a better epistemic situation even if she simply increases her degree of belief in a true hypothesis will *prefer* performing such experiments to maintaining her current degree of confidence in the truth. Unlike a CB model, the PSR model can capture the preferences of such a scientist. For example, the epistemic value of performing \mathbf{X} is almost always positive when the Brier rule is the scientist’s EUF (see Figure 4).




Figure 4. Value of an experiment, using the Brier rule

Even so, it might be suggested that such an attitude toward epistemic risk is unscientific.²⁸ In other words, it might be suggested that scientists should only value believing true hypotheses (and not believing false hypotheses). However, this does not fit with actual scientific practice. As noted above, if scientists only value believing true hypotheses, then the epistemic value of performing an experiment will often be zero. However, if the epistemic value of performing a particular experiment is zero, then it will probably not be rational *all things considered* for scientists to perform it since there are often significant non-epistemic costs to performing an experiment.²⁹

We might try to avoid this consequence by finding another source of epistemic utility. For example, even if a scientist's doxastic state with respect to the hypotheses in H does not change as a result of performing an experiment, she will at least have found out what the result of the experiment was (cf. Levi's [39, pp. 377–378] critique of Goossens [19]). However, since she performed the experiment in order to find out which hypothesis is true, the epistemic utility of simply knowing what the result is is going to be rather small.³⁰ Alternatively, we might try to avoid this consequence by assigning epistemic utilities to a more fine-grained set of doxastic states. For example, we might assign a higher epistemic utility to *strongly* believing a true hypothesis than we assign to *merely* believing that hypothesis. This will increase the number of experiments that have the potential to change scientists' doxastic states. But, since we are already headed in that direction, why not simply assign epistemic utilities directly to degrees of belief in a hypothesis?

Since a scientist's probability assignment is her doxastic state in a DB model, most experiments have the potential to change her doxastic state. As a result, unlike a CB model, the epistemic value of performing an experiment is typically positive in the PSR model (see Figure 4).³¹ Thus, the PSR model fits much better with actual scientific practice.

²⁸This is essentially what Maher's [40, p. 116] rejection of DB models amounts to.

²⁹Even if a particular experiment has no epistemic value, it might be part of a *sequence* of experiments that has epistemic value. However, it seems unlikely that scientists will always have such a sequence in mind.

³⁰Thus, it will probably still not be rational *all things considered* for a scientist to perform the experiment.

³¹My conjecture is that the epistemic value of performing an experiment is positive as long as u_i is a *strictly* increasing function of r_i .

5.2. The Horwich Model

The first DB model of epistemic utilities was proposed by Horwich [24, pp. 127–129]. The very same model of epistemic utilities was proposed independently by Goldman [15, p. 90].³² The Horwich model contains just one EUF which I will refer to as the *linear rule*: $u_i(\mathbf{r}) = r_i$ for all i (see Figure 5).³³




Figure 5. The linear rule

Horwich [24, pp. 127–129] used this model to explain why it is rational for scientists to perform experiments. However, as Maher [40, pp. 112–113] was the first to point out, since the linear rule is not a proper scoring rule, it does not capture a scientifically acceptable attitude toward epistemic risk (cf. [41, pp. 178–179]).³⁴ Since the Horwich model captures only one attitude toward epistemic risk and it is not scientifically acceptable, the PSR model trivially captures all scientifically acceptable attitudes toward epistemic risk

³²Horwich, Maher, and Oddie all argue that performing an experiment is expected *from the perspective of the scientist* to increase epistemic utility. Goldman [15, p. 121], however, uses this model to argue that, under certain conditions, performing an experiment is *objectively* expected to increase epistemic utility (cf. [8]).

³³Horwich [24] and Goldman [15] simply assume that this is the only scientifically acceptable attitude toward epistemic risk. However, there is no obvious reason to think that epistemic utility is necessarily a linear function of degree of belief in the truth. Analogously, Daniel Bernoulli [2] pointed out that utility in general need not be a linear function of wealth.

³⁴A number of epistemologists and philosophers of science have subsequently pointed out that the linear rule is unscientific because it often requires jumping to conclusions (see [49, pp. 264–265], [55, p. 52], [8, pp. 234–236], [1], [7, pp. 100–101], [34, p. 387]). Michael DePaul [7, pp. 108–109] claims that this problem with the Horwich model casts doubt on any form of “truth consequentialism.” However, as noted in section 4, the PSR model avoids this problem.

that are captured by the Horwich model. In addition, the PSR model clearly captures scientifically acceptable attitudes toward epistemic risk that are not captured by the Horwich model.

5.3. The Oddie Model

Oddie (1997, 538) has proposed another DB model of epistemic utilities that is very similar to the PSR model. This model places an even tighter constraint on EUFs than the PSR model: $EU(\mathbf{r}, \mathbf{r}) > EU(\mathbf{q}, \mathbf{r})$ for all $\mathbf{q} \neq \mathbf{r}$.³⁵ In other words, an EUF should be a *strictly proper scoring rule*.

Oddie [46, pp. 539-541] used this model to explain why it is rational for scientists to perform experiments. However, since the Oddie model is clearly a proper subset of the PSR model, the PSR model captures all scientifically acceptable attitudes toward epistemic risk that are captured by the Oddie model. In addition, as I argue below, the PSR model captures scientifically acceptable attitudes toward epistemic risk that are not captured by the Oddie model.

The PSR model captures all of the scientifically acceptable preferences over epistemic lotteries that are captured by CB models (cf. subsection 5.1). However, the Oddie model does not capture any of these preferences over epistemic lotteries, since they cannot be represented by strictly proper scoring rules. In a CB model, while a scientist never prefers simply changing her probability assignment in the absence of new evidence to maintaining her current probability assignment, she will sometimes be indifferent between the two. In other words, there are $\mathbf{q} \neq \mathbf{r}$ such that $EU(\mathbf{q}, \mathbf{r}) = EU(\mathbf{r}, \mathbf{r})$. In particular, the expected epistemic utility of any probability assignment that leaves her in the same doxastic state in a CB model is the same as the expected epistemic utility of her current probability assignment.

Even so, one apparent advantage of the Oddie model is that conservatism seems to be built into the model. In the Oddie model, a scientist is never indifferent between maintaining her current probability assignment and simply changing her probability assignment in the absence of new evidence. She always prefers the first to the second. As I argue below, however, conservatism still must be explicitly added to both the Oddie and PSR models as a further constraint.

In order to decide whether it is rational for a scientist to engage in a particular activity, we have to evaluate the various doxastic states that the scientist could end up in as a result of engaging in this activity (cf. section 3).

³⁵Oddie [46, p. 539] refers to this constraint as *cogency*.

In addition, we have to evaluate them *from the perspective of those new doxastic states themselves* rather than from the perspective of her current doxastic state. For example, in order to decide whether it is rational to perform an experiment, we have to use \mathbf{r}_e to calculate the expected epistemic utility of a scientist's doxastic state if she gets result e . In other words, the expected epistemic utility is $EU(\mathbf{r}_e, \mathbf{r}_e)$ rather than $EU(\mathbf{r}_e, \mathbf{r})$. If we used \mathbf{r} to calculate the expected epistemic utility, then the epistemic value of performing an experiment would never be greater than zero. This is because $EU(\mathbf{r}, \mathbf{r}) \geq EU(\mathbf{r}_e, \mathbf{r})$ for all \mathbf{r}_e in the Oddie and PSR models.

Now, let \mathbf{r}_d be a scientist's new probability assignment if she takes the drug that will make her certain of her favorite scientific hypothesis. In order to decide whether it is rational to take this drug, we have to use \mathbf{r}_d to calculate the expected epistemic utility of her doxastic state if she takes this drug. In the Oddie and PSR models, however, the expected epistemic utility of taking the drug, $EU(\mathbf{r}_d, \mathbf{r}_d)$, is greater than the expected epistemic utility of not taking the drug, $EU(\mathbf{r}, \mathbf{r})$, as long as she is not already certain of this hypothesis.³⁶ In order to conclude that it is not rational to take this drug, we have to explicitly require that "absent any new information you should not change your cognitive state." EUFs by themselves cannot distinguish between (a) the (scientifically acceptable) changes to one's doxastic state that result from performing an experiment and (b) the (scientifically objectionable) changes to one's doxastic state that result, for example, from taking a drug.

Since conservatism must be added as a further constraint, it is sufficient that EUFs simply be consistent with conservatism. In other words, a scientist should never think that another probability assignment is better than her current probability assignment (as the Horwich model often requires). But she does not have to think that her current probability assignment is better than all the other probability assignments (as the Oddie model requires). She just has to think that it is at least as good as all the others (as the PSR model requires).

6. Objections to the Proper Scoring Rule Model

The previous section established that the PSR model is the least restrictive Bayesian model of epistemic utilities. In this section, I defend the PSR model

³⁶With experiments, the $EU(\mathbf{r}_e, \mathbf{r}_e)$ are weighted by how likely she *currently* thinks that each of the possible results are, but we can assume that she is currently certain of the effectiveness of this drug.

against several important objections. First of all, despite the fact that it is the least restrictive Bayesian model, many people worry that the PSR model is still too restrictive. For example, Maher [40, p. 113] claims that “the requirement that cognitive utility functions have the form of proper scoring rules is ad hoc” and that “there is no independent support for the proposition that scientists have cognitive utility functions which are proper scoring rules” (cf. [41, p. 179]). However, for precisely the reasons that Maher [40, pp. 112–113] himself gives, the PSR constraint is a well-motivated constraint on a scientist’s preferences over epistemic lotteries (cf. [46, pp. 535–538], [47, p. 144]). It only rules out EUFs that are clearly unscientific (cf. section 4). In fact, since the PSR model captures a wider range of scientifically acceptable attitudes toward epistemic risk than the CB models that Maher defends, it is the CB models that arguably place an ad hoc restriction on scientists’ epistemic goals.

Ruth Weintraub [55, p. 61] also thinks that the PSR model is too restrictive. She claims that the PSR model “is not even empirically adequate; people’s utility functions vary enormously.” Now, it is no doubt true that there are many scientifically acceptable attitudes toward epistemic risk (cf. section 2). However, as noted in section 4, the PSR model does capture a wide variety of attitudes toward epistemic risk. Admittedly, as I discuss in the next section, there are special contexts in which only one attitude toward epistemic risk is allowed by the PSR model. However, the PSR constraint by itself does not lead to these unreasonably tight restrictions.

Finally, it should be noted that there is another sense in which the PSR model is not overly restrictive. While the PSR model puts constraints on the way in which a scientist’s preferences over epistemic lotteries are *ordered*, it does not put any constraints on the *strength* of those preferences.³⁷ For example, a scientist must prefer assigning a high probability to the true hypothesis to assigning a low probability to the true hypothesis. However, there are no constraints on how much the scientist must prefer the first to the second.

In addition to the worry that the PSR model is too restrictive, it has also been suggested that the PSR constraint is *unnecessary* if we just want to explain why it is rational for scientists to perform experiments. For example, Maher [40, p. 116] argues for the superiority of CB models of epistemic utilities on the grounds that they do not “impose any restrictions on the

³⁷The same sort of point applies to utility functions in general (see, e.g., [33, p. 150]). For example, the fact that $u(x)$ is equal to two times $u(y)$ does not imply that having x is twice as good as having y .

cognitive utility function to ensure that expected cognitive utility cannot be increased by arbitrarily changing probabilities.” However, since all of the EUFs in CB models implicitly satisfy the PSR constraint, the PSR model does not place an additional constraint on EUFs. It is not clear why simply making this constraint explicit should count against the PSR model.

Robert Stalnaker [53, p. 165] also claims that the PSR constraint is unnecessary. In fact, he tries to establish that it is rational for scientists to perform experiments without assuming that EUFs are proper scoring rules. However, it is not clear that this can be done.³⁸ If we do not require that EUFs satisfy the PSR constraint, then there will be EUFs in our model for which the epistemic value of performing some experiments is negative. In fact, it is even possible for an individual to consistently prefer true belief to error and still think that it is sometimes not a good idea to perform experiments. For example, if $u_i(\mathbf{r}) = -e^{-5r_i}$ for all i , then the epistemic value of performing the experiment described in subsection 5.1 is negative when $0.37 < r_i < 0.63$ (cf. [41, pp. 176–177]).³⁹

Stalnaker [53, pp. 166–167] also criticizes the PSR model on the grounds that the PSR constraint is not restrictive enough. And he may very well be correct that the PSR constraint is not the only constraint on scientifically acceptable EUFs.⁴⁰ In other words, being consistent with conservatism may not be all that is required for scientific acceptability. In fact, as I discuss in the next section, there is at least one other plausible constraint on scientifically acceptable EUFs. This is not a problem when we are simply trying to explain why it is rational for scientists to perform experiments, however. For this purpose, it does not matter if the PSR model includes some EUFs that are not scientifically acceptable.⁴¹ The PSR model can provide a complete explanation as long as it does not exclude any EUFs that are scientifically acceptable.

³⁸Stalnaker’s own argument is fallacious. It only goes through if we beg the question by assuming that “investigation is cost free” epistemically as well as non-epistemically.

³⁹The epistemic value of performing this experiment is negative for an even wider range of prior probabilities if an individual is even more averse to epistemic risk (e.g., if $u_i(\mathbf{r}) = -e^{-10r_i}$ for all i). See [33, pp. 159–160] for the standard measure of risk aversion.

⁴⁰James Joyce [30], for example, proposes a different set of constraints on assigning epistemic utilities to degrees of belief. Joyce uses these constraints to argue that a scientist (at least when she seeks the “truth and nothing but the truth”) must have coherent degrees of belief. See Maher [42] for a critique of Joyce’s proposal. In this paper, I simply take the coherence assumption of the Bayesian models for granted.

⁴¹It would just mean that there are some unscientific people for who it is always rational to perform experiments.

7. Seeking the “Truth and Nothing but the Truth”

Since jumping to conclusions is never scientifically acceptable, EUFs must always satisfy the PSR constraint. However, there are special contexts in which EUFs must satisfy additional constraints. In this section, I show that, in one such context, these additional constraints lead to some very interesting restrictions on a scientist’s attitude toward epistemic risk.

As Levi [37, p. 55] suggests, a scientist may sometimes seek the “truth and nothing but the truth.” In this context, her EUF should certainly satisfy the following constraint: u_i should be a monotonically increasing function of r_i . In other words, if the probability that she assigns to the true hypothesis increases, then the epistemic utility of her doxastic state should also increase (cf. [41, p. 176], [55, p. 55]).⁴² More formally, if $r_i > q_i$, then $u_i(\mathbf{r}) \geq u_i(\mathbf{q})$.⁴³

If there are only two hypotheses in H , all EUFs in the PSR model satisfy this *monotonicity* constraint. If there are three or more hypotheses, however, many EUFs in the PSR model do not. This is ultimately because the epistemic utility of a probability assignment might depend, not only on the probability assigned to the true hypothesis, but also on how probabilities are distributed among the various false hypotheses. With the Brier rule (see section 4), for example, the epistemic utility of a probability assignment is higher if probabilities are spread more evenly among the false hypotheses.⁴⁴ For example, if $\mathbf{r} = (0.6, 0.2, 0.2)$ and $\mathbf{s} = (0.6, 0.4, 0)$, then $u_1(\mathbf{r}) > u_1(\mathbf{s})$. Furthermore, the benefit of having probabilities spread more evenly among the false hypotheses can outweigh the cost of assigning a lower probability to the true hypothesis. For example, if $\mathbf{r} = (0.58, 0.21, 0.21)$ and $\mathbf{s} = (0.6, 0.4, 0)$, then $u_1(\mathbf{r}) > u_1(\mathbf{s})$. Thus, if there are three (or more) hypotheses, the Brier rule does not satisfy the monotonicity constraint.

⁴²There are certainly situations where a lower (higher) degree of confidence in the truth has good (bad) *practical* consequences (cf. [45, pp. 69–70]). As a result, Goldman [17, pp. 218–220] has suggested a model that allows epistemic utility to go up as degree of belief in the truth goes down. But such EUFs are clearly unscientific. In addition, there are situations where a lower degree of confidence in the truth today serves our epistemic goals in the long run (cf. [45, p. 69], [11, pp. 322–323]). But the Bayesian models discussed in this paper focus solely on the current epistemic utility of our beliefs about a single set of hypotheses.

⁴³Strictly speaking, I only require here that, if the probability that she assigns to the true hypothesis increases, then the epistemic utility of her doxastic state does not decrease (cf. [41, p. 211]). Requiring u_i to be a strictly increasing function of r_i would rule out all EUFs in CB models.

⁴⁴This is because, for a fixed value of r_i , $\sum_k r_k^2$ is minimized when all $r_{j \neq i}$ have the same value.

All EUFs in the PSR model do, however, satisfy what I will refer to as the *weak monotonicity* constraint (see Proposition 1 in the Appendix). If the probability assigned to the true hypothesis goes up *and* the probabilities assigned to the false hypotheses are in the same ratios, then the epistemic utility of the probability assignment does not go down. More formally, if $r_i > q_i$ and there is a real number α such that $\alpha r_j = q_j$ for all $j \neq i$, then $u_i(\mathbf{r}) \geq u_i(\mathbf{q})$. The weak monotonicity constraint basically insures that, all other things being equal, if the probability that a scientist assigns to the true hypothesis increases, then the epistemic utility of her doxastic state increases.

It might be suggested that *scientifically acceptable* EUFs must always satisfy the full monotonicity constraint as well as the PSR constraint. Admittedly, if the probability that she assigns to the true hypothesis increases, the probability that she assigns to a simpler or more explanatory hypothesis might decrease. But it is arguable that the simplicity and explanatory power of a hypothesis should only be valuable to her if the hypothesis is true. Thus, if the probability that she assigns to the true hypothesis increases, then the epistemic utility of her doxastic state should still increase.

However, it is also arguable that scientifically acceptable EUFs must only satisfy the weak monotonicity constraint. The epistemic utility of a probability assignment might very well depend on how probabilities are distributed among the various false hypotheses. In particular, scientists might not want all the probability assigned to false hypotheses concentrated on a single false hypothesis. In other words, it might be epistemically better to assign a slightly lower probability to the true hypothesis if probabilities are spread more evenly among the false hypotheses. This seems especially reasonable when the probability concentrated on the single false hypothesis would be greater than the probability assigned to the true hypothesis. For example, it seems quite reasonable to follow the Brier rule and say that $u_1(\mathbf{r}) > u_1(\mathbf{s})$ when $\mathbf{r} = (0.38, 0.31, 0.31)$ and $\mathbf{s} = (0.4, 0.6, 0)$. Thus, I will only assume here that a scientist's EUF must satisfy the full monotonicity constraint when she seeks the "truth and nothing but the truth."

But if a scientist seeks the "truth and nothing but the truth," her EUF should satisfy at least one other constraint in addition to the monotonicity constraint. While the epistemic utility of her probability assignment certainly depends on the probability assigned to the true hypothesis, it should not depend on *which* hypothesis happens to be true (cf. subsection 5.1). More formally, an EUF satisfies the *invariance* constraint if $u_i(\mathbf{r}) = u_j(\mathbf{r}^{(i,j)})$ for all i and j and \mathbf{r} where $\mathbf{r}^{(i,j)}$ is the same as \mathbf{r} with r_i and r_j swapped (cf. [52, p. 38]).

If there are two mutually exclusive and jointly exhaustive hypotheses in H , there are many different EUFs that satisfy the monotonicity and invariance constraints as well as the PSR constraint. For example, all of the EUFs in Levi's original model satisfy all three of these constraints. However, if there are three or more hypotheses in H , the logarithmic rule is the only continuous EUF that satisfies all three constraints (see Proposition 4 in the Appendix).⁴⁵ Thus, when a scientist seeks the "truth and nothing but the truth" and there are three or more hypotheses, there is essentially only one scientifically acceptable attitude toward epistemic risk (see Figure 6).⁴⁶

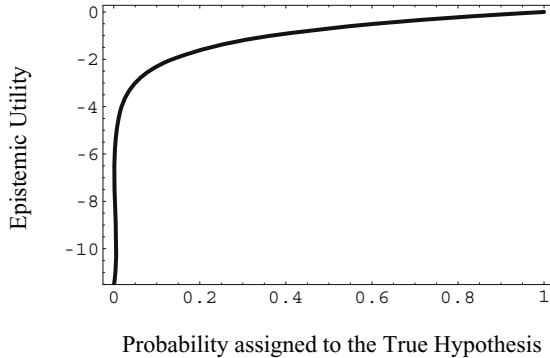


Figure 6. The logarithmic rule

It should be noted that, as it stands, this result only shows that the logarithmic rule is the only *continuous* function that satisfies the three constraints. It does not, for example, rule out the possibility that there are EUFs in CB models that satisfy the three constraints. As noted in subsection 5.1, EUFs in CB models are (discontinuous) step functions when they are translated into a DB model (see Figure 2).

However, there are no step functions over three or more hypotheses that satisfy the PSR, invariance, and monotonicity constraints (see Proposition

⁴⁵This result is based on a theorem that was originally proved by Shuford, Albert, and Massengill [50, pp. 135–137]. For a concise proof of this theorem applied to the case of epistemic utilities, see [9, p. 180]. Finally, it should be noted that, strictly speaking, the logarithmic rule itself is discontinuous, but only at $r_i = 0$.

⁴⁶Even if a scientist is not seeking the "truth and nothing but the truth," the invariance constraint still must be satisfied by her EUF whenever all of the hypotheses in H are equivalent with respect to her other epistemic goals (e.g., simplicity or explanatory power). Thus, if it does turn out that scientifically acceptable EUFs must always satisfy the full monotonicity constraint, there will be many more situations where the logarithmic rule is the only scientifically acceptable attitude toward epistemic risk.

6 in the Appendix).⁴⁷ As a result, if an EUF in a CB model satisfies the invariance constraint, then it does not satisfy the monotonicity constraint. That is, epistemic utility will sometimes go down even if the scientist's degree of belief in the truth goes up. For example, Levi [38, p. 81] proposed an EUF over three hypotheses that satisfies the invariance constraint. If we let the scientist's "degree of caution" (q) be 0.75, then $u_1(\mathbf{r}) > u_1(\mathbf{s})$ when $\mathbf{r} = (0.58, 0.21, 0.21)$, and $\mathbf{s} = (0.6, 0.4, 0)$.

My conjecture is that there are no discontinuous functions that satisfy the PSR, invariance, and monotonicity constraints. However, the EUFs in CB models are the only discontinuous functions that are reasonable candidates for a scientist's attitude toward epistemic risk.⁴⁸ Thus, even without a proof of this conjecture, we are left with the conclusion that the logarithmic rule is the only scientifically acceptable attitude toward epistemic risk when a scientist seeks the "truth and nothing but the truth."

8. Unreasonably Tight Restrictions

The logarithmic rule is almost certainly a scientifically acceptable EUF. In fact, as I discuss below, it has several very attractive properties.⁴⁹ However, it is not clear why a scientist who seeks the "truth and nothing but the truth" should be *required* to have this particular attitude toward epistemic risk.⁵⁰ In other words, Bayesian models of epistemic utilities seem to place unreasonably tight restrictions on a scientist's attitude toward epistemic risk in this context.

⁴⁷The proof shows that, if a step function is invariant and monotonic, then it is not a proper scoring rule. This is done by constructing a probability assignment that "thinks" that another probability assignment is epistemically better. This second probability assignment is essentially the same as the first probability assignment except that it assigns enough additional probability to one hypothesis to land on the next higher step of the step function. (But the construction of these probability assignments is a little tricky. In order to increase the probability assigned to one hypothesis, we have to decrease the probability assigned to another hypothesis. And we have to do this in such a way that the probability assigned to this other hypothesis remains on the same step of the step function.)

⁴⁸As far as I know, no one has ever proposed any other discontinuous functions as EUFs.

⁴⁹Interestingly, while Goldman's claim [15, p. 121] that performing an experiment is *objectively* expected to increase epistemic utility does not hold for any bounded EUF that satisfies the PSR constraint, it does hold for the logarithmic rule (cf. [9, p. 178]).

⁵⁰Bernoulli [2, p. 28] assumed that utility in general was a logarithmic function of wealth. However, as Karl Menger [44, p. 228] pointed out, this was an "ad hoc" assumption. There are many other possible utility functions. Similarly, assuming that epistemic utility is a logarithmic function of degree of belief in the truth seems "ad hoc."

First, with the logarithmic rule, the cost of decreasing one's degree of belief in the truth by a particular amount is always greater than the benefit of increasing one's degree of belief in the truth by that same amount. Thus, if a scientist's attitude toward epistemic risk is captured by the logarithmic rule, she is averse to epistemic risk (cf. [55, p. 61]).⁵¹ This seems perfectly reasonable. Most epistemologists (e.g., Descartes and Hume) claim that people should be cautious in their inquiries (cf. [8, p. 232], [48, p. 347], [7, p. 98]). However, it also seems reasonable for a scientist to be more or less cautious than the logarithmic rule requires.⁵²

Second, with the logarithmic rule, as one's degree of belief in the truth goes up, the cost of decreasing one's degree of belief in the truth by a particular amount goes down relative to the benefit of increasing one's degree of belief in the truth by that same amount. Thus, if a scientist's attitude toward epistemic risk is captured by the logarithmic rule, she is more willing to take epistemic risks as she gets closer to the truth (cf. [33, p. 169]). This seems reasonable. People's utility functions often exhibit this sort of *decreasing risk aversion* since "as they become richer, they can better afford to take a specific risk" [33, p. 166]. However, it also seems reasonable for a scientist to be less willing to risk making a mistake as she gets closer to the truth.⁵³

Third, with the logarithmic rule, the epistemic utility of having a degree of belief of zero in the true hypothesis is $-\infty$. Thus, if a scientist's attitude toward epistemic risk is captured by the logarithmic rule, then she assigns infinite disutility to getting things completely wrong.⁵⁴ This seems reasonable since getting things completely wrong is pretty bad.⁵⁵ If a scientist has a degree of belief of zero in a hypothesis and she updates her probability assignment via conditionalization, no amount of evidence will ever lead her

⁵¹An individual is risk averse if her utility function is concave (cf. [33, p. 149]).

⁵²In a similar vein, Maher [40, p. 113] suggests that she might reasonably be perfectly risk neutral (i.e., that the linear rule might capture her attitude toward epistemic risk).

⁵³The Brier rule exhibits this sort of increasing risk aversion (cf. [33, p. 168]).

⁵⁴Assigning infinite disutility to something does not mean that it is *infinitely bad*. A utility function only captures the way in which an individual's preferences are *ordered* and not the *strength* of those preferences (cf. section 6). In particular, assigning infinite disutility to believing a false hypothesis just means that any epistemic lottery that gives a positive probability to believing a false hypothesis is at the very bottom of the scientist's preference ranking.

⁵⁵Similarly, it seems reasonable to assign infinite disutility to damnation when faced with Pascal's wager (cf. [6, p. 510]). Of course, a scientist will never assign positive probability to being completely wrong as a result of performing an experiment. As a result, scientists may never actually be confronted with such epistemic lotteries.

to increase her degree of belief in that hypothesis (cf. [55, pp. 61-62]). In other words, she is doomed to get things completely wrong forever. However, it also seems reasonable for a scientist to have some other attitude toward getting things completely wrong. In fact, *boundedness* is a fairly standard constraint on utility functions (see, e.g., [6, p. 505]).⁵⁶

Of course, the Bayesian can always bite the bullet. In other words, she might accept that, strangely enough, the logarithmic rule is the only scientifically acceptable attitude toward risk when seeking the “truth and nothing but the truth” since this result follows from the Bayesian models of epistemic utilities. In addition, there are at least a couple of other responses that the Bayesian might make to the logarithm result. But, however the Bayesian responds, there turn out to be unreasonably tight restrictions on a scientist’s epistemic goals.

First, the Bayesian might retreat to a CB model of epistemic utilities (as Maher recommends). Now, as noted above, there are no EUFs in CB models that satisfy the PSR, invariance, and monotonicity constraints. But, as noted in subsection 5.1, all of the EUFs in CB models do satisfy a restricted monotonicity constraint. That is, it is always epistemically better to believe the true hypothesis than to withhold judgment, and it is always epistemically better to withhold judgment than to believe a false hypothesis. However, in retreating to a CB model, we have to give up the perfectly reasonable idea that it is epistemically better to simply increase one’s degree of belief in the true hypothesis (cf. subsection 5.1). Thus, even if the Bayesian responds in this way, there are unreasonably tight restrictions on a scientist’s epistemic goals when she seeks the “truth and nothing but the truth.”

Second, the Bayesian might claim that scientists should never seek the “truth and nothing but the truth.” For example, she might claim that scientists should always be concerned with simplicity and explanatory power as well as with the truth. However, taking this position arguably places unreasonably tight restrictions on a scientist’s epistemic goals. It is not immediately clear why it would be *unscientific* for a scientist to seek the “truth and nothing but the truth.” In fact, many philosophers (e.g., Goldman [16]) claim that that is all that scientists ever do. They argue that any epistemic goals beyond true belief, such as simplicity or explanatory power, are only valuable as a means to true belief. Thus, even if scientists are always concerned with simplicity and explanatory power, they are still seeking the “truth and nothing but the truth.”

⁵⁶Weintraub [55, p. 54] actually claims that EUFs should be bounded, but she offers no argument for this claim.

9. Conclusion

The PSR model of epistemic utilities places a single, well-motivated constraint on scientists' attitudes toward epistemic risk. In particular, a scientist should not be able to increase her expected epistemic utility by simply changing her probability assignment in the absence of new evidence. As a result of this, the PSR model captures a wider range of scientifically acceptable attitudes toward epistemic risk than the other Bayesian models of epistemic utilities. Thus, this Bayesian model of epistemic utilities provides the most comprehensive explanation for why it is rational for scientists to perform experiments. Unfortunately, any such Bayesian model places unreasonably tight restrictions on a scientist's attitude toward epistemic risk when she seeks the "truth and nothing but the truth." As a result, Bayesian models of epistemic utilities fail as normative accounts of scientific inquiry.

10. Appendix

10.1. Definitions

- Monotonicity: If $r_i > q_i$, then $u_i(\mathbf{r}) \geq u_i(\mathbf{q})$.
- Weak Monotonicity: If $r_i > q_i$, and there is a real number α such that $\alpha r_j = q_j$ for all $j \neq i$, then $u_i(\mathbf{r}) \geq u_i(\mathbf{q})$.
- Relevance: If $r_i = q_i$, then $u_i(\mathbf{r}) = u_i(\mathbf{q})$.
- Invariance: $u_i(\mathbf{r}) = u_j(\mathbf{r}^{(i,j)})$ for all i and j and \mathbf{r} where $\mathbf{r}^{(i,j)}$ is the same as \mathbf{r} with r_i and r_j swapped.
- Proper Scoring Rule (PSR): $\sum_k r_k u_k(\mathbf{r}) \geq \sum_k r_k u_k(\mathbf{q})$ for all \mathbf{q} .

10.2. Propositions

PROPOSITION 1. *If an EUF satisfies the PSR constraint, then it also satisfies the weak monotonicity constraint.*

PROOF. Suppose that an EUF (i.e., a family of u_i functions) satisfies the PSR constraint, but does not satisfy the weak monotonicity constraint. In that case, there must be an i and two probability assignments, \mathbf{r} and \mathbf{q} , such that:

- (a) $r_i > q_i$,
- (b) there is a real number α such that $\alpha r_k = q_k$ for all $k \neq i$,
- (c) but $u_i(\mathbf{r}) < u_i(\mathbf{q})$.

Let $\sum_k^{-i} r_k u_k(\mathbf{r}) = \sum_k r_k u_k(\mathbf{r}) - r_i u_i(\mathbf{r})$. In other words, $\sum_k^{-i} r_k u_k(\mathbf{r})$ is the sum of the $r_k u_k(\mathbf{r})$ with k ranging from 1 to n , but with $r_i u_i(\mathbf{r})$ left out.

Since this EUF satisfies the PSR constraint, $\sum_k r_k u_k(\mathbf{r}) \geq \sum_k r_k u_k(\mathbf{q})$. Thus, $\sum_k^{-i} r_k u_k(\mathbf{r}) + r_i u_i(\mathbf{r}) \geq \sum_k^{-i} r_k u_k(\mathbf{q}) + r_i u_i(\mathbf{q})$. Thus, $\sum_k^{-i} r_k u_k(\mathbf{r}) - \sum_k^{-i} r_k u_k(\mathbf{q}) \geq r_i u_i(\mathbf{q}) - r_i u_i(\mathbf{r}) = r_i (u_i(\mathbf{q}) - u_i(\mathbf{r}))$. Call this fact 1.

Since this EUF satisfies the PSR constraint, $\sum_k q_k u_k(\mathbf{q}) \geq \sum_k q_k u_k(\mathbf{r})$. Thus, $\sum_k^{-i} q_k u_k(\mathbf{q}) + q_i u_i(\mathbf{q}) \geq \sum_k^{-i} q_k u_k(\mathbf{r}) + q_i u_i(\mathbf{r})$. Thus, $q_i (u_i(\mathbf{q}) - u_i(\mathbf{r})) = q_i u_i(\mathbf{q}) - q_i u_i(\mathbf{r}) \geq \sum_k^{-i} q_k u_k(\mathbf{r}) - \sum_k^{-i} q_k u_k(\mathbf{q})$. Call this fact 2.

From (b), $\sum_k^{-i} r_k u_k(\mathbf{r}) - \sum_k^{-i} q_k u_k(\mathbf{q}) = \sum_k^{-i} \alpha r_k u_k(\mathbf{r}) - \sum_k^{-i} \alpha r_k u_k(\mathbf{q}) = \alpha (\sum_k^{-i} r_k u_k(\mathbf{r}) - \sum_k^{-i} r_k u_k(\mathbf{q}))$. From (a) and the fact that \mathbf{r} and \mathbf{q} are probability assignments, $\alpha > 1$. From (c), the fact that \mathbf{r} is a probability assignment, and fact 1, $\sum_k^{-i} r_k u_k(\mathbf{r}) - \sum_k^{-i} r_k u_k(\mathbf{q}) \geq 0$. Thus, $\sum_k^{-i} q_k u_k(\mathbf{r}) - \sum_k^{-i} q_k u_k(\mathbf{q}) \geq \sum_k^{-i} r_k u_k(\mathbf{r}) - \sum_k^{-i} r_k u_k(\mathbf{q})$. Call this fact 3.

Putting together facts 1 through 3, we find that $q_i (u_i(\mathbf{q}) - u_i(\mathbf{r})) \geq r_i (u_i(\mathbf{q}) - u_i(\mathbf{r}))$. From (c), it follows that $q_i \geq r_i$, but this contradicts (a). ■

THEOREM 2 (Shuford, Albert, and Massengill 1966). *When there are $n \geq 3$ hypotheses in H , $u_i(\mathbf{r}) = \log(r_i)$ for all i is the only continuous EUF that satisfies the relevance, invariance, and PSR constraints.*

LEMMA 3. *If a continuous EUF satisfies the monotonicity constraint, then it also satisfies the relevance constraint.*

PROOF. Suppose that a continuous EUF (i.e., a family of continuous u_i functions) satisfies the monotonicity constraint, but does not satisfy the relevance constraint. In that case, there must be an i and two probability assignments, \mathbf{r} and \mathbf{q} , such that:

- $r_i = q_i$,
- but $u_i(\mathbf{r}) \neq u_i(\mathbf{q})$.

Without loss of generality, suppose that $u_i(\mathbf{r}) > u_i(\mathbf{q})$. Let $\varphi = u_i(\mathbf{r}) - u_i(\mathbf{q})$.

Since u_i is continuous, there must be a probability assignment \mathbf{s} such that $s_i = q_i + \delta$, $\delta > 0$, and $u_i(\mathbf{s}) - u_i(\mathbf{q}) < \varphi$. Thus, $u_i(\mathbf{r}) > u_i(\mathbf{s})$.

But since u_i is monotonic and $s_i > r_i$, $u_i(\mathbf{s}) \geq u_i(\mathbf{r})$. This is a contradiction. ■

PROPOSITION 4. *When there are $n \geq 3$ hypotheses in H , $u_i(\mathbf{r}) = \log(r_i)$ for all i is the only continuous EUF that satisfies the monotonicity, invariance, and PSR constraints.*

PROOF. This result follows trivially from theorem 2 and lemma 3. ■

LEMMA 5. *If an EUF made up of step functions with finitely many steps satisfies the monotonicity constraint, then it also satisfies the relevance constraint-except possibly at the locations of the discontinuities of the step functions.*

PROOF. The proof is essentially the same as the proof of lemma 3. We simply pick r_i and q_i such that they are not at the location of a discontinuity. ■

PROPOSITION 6. *When there are $n \geq 3$ hypotheses in H , there is no EUF made up of step functions with finitely many steps that satisfies the monotonicity, invariance, and PSR constraints.*

PROOF. Suppose that an EUF (i.e., a family of u_i functions) made up of step functions with finitely many steps satisfies the monotonicity and invariance constraints. Since this EUF satisfies the invariance constraint, the u_i functions are essentially the same function. Thus, the locations of the discontinuities are the same for all of the u_i functions.

Let δ be the length of the shortest step. Let γ be the location of the first discontinuity.

Note: Let \mathbf{s} be such that $u_1(\mathbf{s})$ is the minimum epistemic utility. The location of the first discontinuity is the smallest value for r_1 such that $u_1(\mathbf{r}) > u_1(\mathbf{s})$. Since the relevance constraint is not necessarily satisfied at the discontinuities, the actual value of $u_1(\mathbf{r})$ might depend on the values of the other r_k as well as the value of r_1 . However, since (by lemma 5) the relevance constraint is satisfied everywhere else, the locations of the discontinuities only depend on the value of r_1 .

First, construct a probability assignment, \mathbf{r} , such that:

- r_1 is right before (viz., within 0.1δ of) the next step,
- r_2 is far past (viz., at least 0.5δ beyond) the previous step,
- and r_k is somewhere in the middle of a step for all k .

For example, let $\varphi = 1 - (\gamma - 0.1\delta) - 0.5\delta$. If $\varphi/(n - 2) < \gamma$, then let:

- $r_1 = \gamma - 0.1\delta$
- $r_2 = 0.5\delta$,
- and $r_k = \varphi/(n - 2)$ for all $k \geq 3$.

Note: In this case, all of the r_k come before the first step.

Otherwise, let $\mu = 1 - (n - 2)(\gamma - 0.1\delta) - 0.5\delta$. If μ is not the location of a discontinuity, then let:

- $r_1 = \gamma - 0.1\delta$
- $r_2 = 0.5\delta$,
- $r_k = \gamma - 0.1\delta$ for all $k \geq 3$ except n ,
- and $r_n = \mu$.

Note: In this case, all of the r_k except r_n come before the first step.

Otherwise, let:

- $r_1 = \gamma - 0.1\delta$
- $r_2 = 0.6\delta$,
- $r_k = \gamma - 0.1\delta$ for all $k \geq 3$ except n ,
- and $r_n = \mu - 0.1\delta$.

Note: This is just like the previous case except that r_2 and r_n are adjusted slightly so that r_n is not the location of a discontinuity.

Next, construct a probability assignment, \mathbf{q} , such that:

- $q_1 = r_1 + 0.2\delta$,
- $q_2 = r_2 - 0.2\delta$,
- and $q_k = r_k$ for all $k \geq 3$.

Since q_k and r_k are somewhere in the middle of a step for all k , $u_k(\mathbf{q})$ is (by lemma 5) a function of q_k and $u_k(\mathbf{r})$ is a function of r_k for all k . This allows us to make the following three inferences:

(1) Since r_1 is right before the next step and q_1 is right after this step and u_1 is monotonic, $u_1(\mathbf{q}) > u_1(\mathbf{r})$. In other words, when we add 0.2δ to r_1 , epistemic utility goes up because we move up to a higher step.

(2) Since q_2 and r_2 are on the same step, $u_2(\mathbf{q}) = u_2(\mathbf{r})$.

(3) Since $q_k = r_k$ for all $k \geq 3$, $u_k(\mathbf{q}) = u_k(\mathbf{r})$ for all $k \geq 3$.

It follows from these three facts that $\sum_k r_k u_k(\mathbf{q}) > \sum_k r_k u_k(\mathbf{r})$. Thus, this EUF does not satisfy the PSR constraint. ■

Acknowledgements. I would like to thank Brad Armendt, Jeff Barrett, David Christensen, Andrew Cohen, Martin Frické, Alvin Goldman, Risto Hilpinen, Peter Lewis (especially Pete!), Isaac Levi, Gerrard Liddell, Kay Mathiesen, Brian Skyrms, Paul Weirich, and Dennis Whitcomb for many helpful comments and suggestions on earlier versions of this paper.

References

- [1] BERENDS, H., 'Veritistic value and the use of evidence: A shortcoming of Goldman's epistemic evaluation of social practices', *Social Epistemology* 16:177–79, 2002.
- [2] BERNOULLI, D., 'Exposition of a new theory on the measurement of risk', *Econometrica* 22:23–36, 1954.
- [3] BORGMAN, C. L., and J. FURNER, 'Scholarly communication and bibliometrics', in B. CRONIN (ed.), *Annual Review of Information Science and Technology*, vol. 36, Information Today, 2002, pp. 3–72.
- [4] BOVENS, L., and S. HARTMANN, *Bayesian Epistemology*, Oxford, 2003.
- [5] CHRISTENSEN, D., *Putting Logic in Its Place*, Oxford, 2004.
- [6] CLEMEN, R. T., *Making Hard Decisions*. 2nd ed., Duxbury Press, 1996.
- [7] DEPAUL, M. R., 'Truth consequentialism, withholding and proportioning belief to the evidence', *Philosophical Issues* 14:91–112, 2004.
- [8] FALLIS, D., 'Goldman on probabilistic inference', *Philosophical Studies* 109:223–40, 2002.
- [9] FALLIS, D., and G. LIDDELL, 'Further results on inquiry and truth possession', *Statistics and Probability Letters* 60:169–82, 2002.
- [10] FIELD, H., 'Realism and relativism', *Journal of Philosophy* 79:553–67, 1982.
- [11] FIRTH, R., *In Defense of Radical Empiricism*, Rowman and Littlefield, 1998.
- [12] FOLEY, R., 'The epistemology of belief and the epistemology of degrees of belief', *American Philosophical Quarterly* 29:111–24, 1992.
- [13] FRANKLIN, A., and C. HOWSON, 'It probably is a valid experimental result: A Bayesian approach to the epistemology of experiment', *Studies in History and Philosophy of Science* 19:419–27, 1988.
- [14] GIERE, R. N., *Explaining Science*, University of Chicago, 1988.
- [15] GOLDMAN, A. I., *Knowledge in a Social World*, Oxford, 1999.
- [16] GOLDMAN, A. I., 'The unity of the epistemic virtues', in A. Fairweather, and L. Zagzebski (eds.), *Virtue Epistemology*, Oxford, 2001, pp. 30–48.
- [17] GOLDMAN, A. I., 'Reply to commentators', *Philosophy and Phenomenological Research* 64:215–27, 2002.
- [18] GOOD, I. J., 'On the principle of total evidence', *British Journal for the Philosophy of Science* 17:319–22, 1967.
- [19] GOOSENS, W. K., 'A critique of epistemic utilities', in R. J. Bogdan, D. Reidel (ed.), *Local Induction*, 1976, pp. 93–113.
- [20] HARDING, S., 'After the neutrality ideal: Science, politics, and "strong objectivity"', *Social Research* 59:567–87, 1992.
- [21] HAREL, D., *The Science of Computing*, Addison-Wesley, 1989.
- [22] HAWTHORNE, J., and L. BOVENS, 'The preface, the lottery, and the logic of belief', *Mind* 108:241–64, 1999.
- [23] HEMPEL, C. G., 'Inductive inconsistencies', *Synthese* 12:439–69, 1960.
- [24] HORWICH, P., *Probability and Evidence*, Cambridge, 1982.
- [25] HOWSON, C., and P. URBACH, *Scientific Reasoning*, Open Court, 1989.
- [26] HUME, D., in E. Steinberg (ed.), *An Enquiry Concerning Human Understanding*, Hackett, 1977.

- [27] HUNTER, D., 'On the relation between categorical and probabilistic belief', *Nous* 30:75–98, 1996.
- [28] JAMES, W., *The Will to Believe and Other Essays in Popular Philosophy*, Harvard, 1979.
- [29] JEFFREY, R., 'Valuation and acceptance of scientific hypotheses', *Philosophy of Science* 23:237–46, 1956.
- [30] JOYCE, J. M., 'A nonpragmatic vindication of probabilism', *Philosophy of Science* 65:575–603, 1998.
- [31] KAHNEMAN, D., and A. TVERSKY, 'Subjective probability: A judgment of representativeness', *Cognitive Psychology* 3:430–454, 1972.
- [32] KAPLAN, M., 'Decision theory and epistemology', in P. K. Moser (ed.), *The Oxford Handbook of Epistemology*, Oxford, 2002, pp. 434–62.
- [33] KEENEY, R. L., and H. RAIFFA, *Decisions With Multiple Objectives*, Cambridge, 1993.
- [34] KIERLAND, B., and B. MONTON, 'Minimizing inaccuracy for self-locating beliefs', *Philosophy and Phenomenological Research* 70:384–95, 2005.
- [35] KITCHER, P., 'Knowledge and tradition', *Philosophical Topics* 29:251–70, 2001.
- [36] LEHRER, K., 'Reason and consistency', in K. Lehrer (ed.), *Analysis and Metaphysics*, Reidel, 1975, pp. 57–74.
- [37] LEVI, I., 'On the seriousness of mistakes', *Philosophy of Science* 29:47–65, 1962.
- [38] LEVI, I., *Gambling With Truth*, MIT, 1967.
- [39] LEVI, I., 'Epistemic utility and the evaluation of experiments', *Philosophy of Science* 44:368–86, 1977.
- [40] MAHER, P., 'Why scientists gather evidence', *British Journal for the Philosophy of Science* 41:103–19, 1990.
- [41] MAHER, P., *Betting on Theories*, Cambridge, 1993.
- [42] MAHER, P., 'Joyce's argument for probabilism', *Philosophy of Science* 69:73–81, 2002.
- [43] MCCLOSKEY, H. J., 'An examination of restricted utilitarianism', *Philosophical Review* 66:466–85, 1957.
- [44] MENGER, K., 'The role of uncertainty in economics', in M. Shubik (ed.), *Essays in Mathematical Economics*, Princeton, 1967, pp. 211–31.
- [45] NOZICK, R., *The Nature of Rationality*, Princeton, 1993.
- [46] ODDIE, G., 'Conditionalization, cogency, and cognitive value', *British Journal for the Philosophy of Science* 48:533–41, 1997.
- [47] PERCIVAL, P., 'Epistemic consequentialism', *Supplementary Volume: Aristotelian Society* 76:121–51, 2002.
- [48] RIGGS, W. D., 'Balancing our epistemic ends', *Nous* 37:342–52, 2003.
- [49] SCHMITT, F. F., 'Veritistic value', *Social Epistemology* 14:259–80, 2000.
- [50] SHUFORD, E. H., A. ALBERT, and H. E. MASSENGILL, 'Admissible probability measurement procedures', *Psychometrika* 31:125–45, 1966.
- [51] SKYRMS, B., *The Dynamics of Rational Deliberation*, Harvard, 1990.
- [52] STAËL VON HOLSTEIN, C.-A. S., *Assessment and Evaluation of Subjective Probability Distributions*, Economic Research Institute, 1970.
- [53] STALNAKER, R., 'Epistemic consequentialism', *Supplementary Volume: Aristotelian Society* 76:153–68, 2002.

- [54] WEATHERSON, B., 'Can we do without pragmatic encroachment?', *Philosophical Perspectives* 19:417–43, 2005.
- [55] WEINTRAUB, R., 'A Bayesian paradox', *British Journal for the Philosophy of Science* 52:51–66, 2001.

DON FALLIS
School of Information Resources
University of Arizona
1515 East First Street
Tucson, Arizona, USA
fallis@email.arizona.edu