

Epistemic Values and the Value of Learning

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario

Abstract. In addition to purely practical values, there are *cognitive values* which figure in scientific deliberations. One way of introducing cognitive values is to consider the cognitive value that accrues to the act of *accepting* a hypothesis. Although such values may have a role to play in the matter of theory acceptance, this does not exhaust their significance in scientific decision-making. This paper makes a plea for the consideration of *epistemic* value—cognitive value that attaches to a state of *belief*. I defend the notion of cognitive epistemic value against criticisms that have been raised against it. A stability requirement for epistemic value-functions is argued for on the basis of considerations of diachronic coherence. This requirement is sufficient for proving the Value of Learning Theorem, which says that the expected utility of cost-free learning cannot be negative. Under the assumption of stability, the expected cognitive epistemic value of undergoing a learning experience must also be non-negative.

1. Varieties of Theory Choice

Discussions of theory choice in science frequently cite theoretical, or cognitive virtues as relevant to theory assessment. Kuhn's well-known list includes accuracy, consistency, scope, simplicity, and fruitfulness (Kuhn, 1977). Explanatory capacity is also commonly cited as a theoretical virtue.

Virtues such as these are desirable features; we would like our theories to have them. Having said that, the question arises as to what, if any, legitimate roles these virtues may play in scientific deliberations. If we regard theories that have these desirable features as, *ceteris paribus*, more credible than theories that lack them, this smacks of wishful thinking. If we reject this, is there any legitimate role for them to play at all?

The phrase “theory choice,” which often occurs in such discussions, obscures the rather obvious and commonplace fact that there are a number of distinct choices that we can make with regards to theories. Compare, for instance, choosing which theory one regards as most credible (or most likely to be either true or approximately true), with the choice of which theory to devote research effort to. These are distinct choices, to which different considerations apply. Credibility of a theory is not the only consideration relevant to the question of choice of research programme; one might even decide to devote one's efforts to what one regards as a ‘long shot,’ if the possible benefits, tangible or cognitive, that would accrue should the work pan out are high enough.

Discussions of theory choice vary in explicitness about what sorts of decisions are involved, and, as a consequence, lists of criteria for theory choice sometimes include criteria relevant to different sorts of choice. This definitely seems to be the case for Kuhn's list.¹ This can

¹ This is also the case for Longino's proposed supplementation to that list (Longino, 1995). Longino is explicit that her list combines criteria relevant to at least three distinct choices: theory construction, theory acceptance or rejection, and (borrowing a term from

result in it not being clear whether a given author is making a claim that cognitive values play a legitimate role in assessing the credibility of a theory—a claim that would rightly be regarded as controversial—or the much less controversial claim that they may play a legitimate role in other decisions, such as the decision of what line of research to pursue.

One decision a scientist may make is the decision whether to *accept* a hypothesis. This, it seems, is a decision that ineliminably involves values, as was famously argued long ago by Rudner (1953). One way to make explicit the role of values in acceptance-decisions is to model the decision as one of maximizing expected utility. In addition to considerations of practical gain or cost, *cognitive* utility may be weighed in the balance. The idea is that knowledge can be regarded as an end in itself, and that acceptance of a true theory may have value that extends beyond instrumental value. This opens up the possibility that acceptance of a theory that has some virtue—*e.g.* explanatory power—might be more valuable than acceptance of one that lacks that virtue. That is, if h_1 is a theory that, if true, affords greater understanding than h_2 would, if true, then one might regard acceptance of h_1 , if it is true, as having higher cognitive utility than acceptance of h_2 if it is true. This, in turn, will be relevant to choice of research programme; it may make sense to investigate theories of high explanatory power, not because one regards them as inherently more credible, but because of the greater cognitive utility that would be realized should the investigation lead to acceptance of the theory.²

Denote by A_h the act of accepting a hypothesis h , and let $u(X|h)$ be the utility of doing X if h is true, $u(X|\sim h)$, the utility of doing X if h is false. Then the expected utility of accepting h , as judged by an agent whose utilities are represented by u and whose credences are represented by a credence-function cr , is

$$U(A_h) = cr(h) u(A_h|h) + cr(\sim h) u(A_h|\sim h).$$

The expected utility of not accepting h is

$$U(\sim A_h) = cr(h) u(\sim A_h|h) + cr(\sim h) u(\sim A_h|\sim h).$$

An agent who maximizes utility will accept h if $U(A_h) > U(\sim A_h)$. This makes the threshold for accepting h —that is, how high $cr(h)$ has to be in order to warrant acceptance of h —depend explicitly on utilities; accepting h is favoured over not accepting it if and only if

$$\frac{cr(h)}{cr(\sim h)} > \frac{u(\sim A_h|\sim h) - u(A_h|\sim h)}{u(A_h|h) - u(\sim A_h|h)}.$$

Alan Franklin), theory pursuit, the choice of which theories or theoretical frameworks to work on.

² A note on terminology. We will use the term *cognitive utility*, or *cognitive value*, for value associated with cognitive goals such as knowledge or understanding, not stemming from the consequences of practical decisions. Cognitive utilities may attach to the act of acceptance, or to a belief state. Following Greaves and Wallace (2006), we will use the term *epistemic utility* for a value attached to a belief state. A belief state may be evaluated in terms of its effect on decisions, in which case we are concerned with *practical* epistemic utilities. It is one goal of this paper to convince the reader that it is both permissible and worthwhile to consider also cognitive epistemic utilities. Levi (1967) and Hempel (1981) use the phrase “epistemic utility” for what we will call the cognitive utility attached to acceptance. Sympathy is extended to any reader who finds the terminological variance in the literature confusing.

An acceptance-based approach to cognitive utility has been developed by a number of authors, notably Hempel (1965; 1981), Levi (1967), Maher (1993), and Douglas (2000). There is much to be said for such an approach. In particular, it makes room for the relevance of cognitive values to decisions regarding theory pursuit. Note that, though cognitive values may play a role in a decision to accept a theory or to pursue it, these values need not affect the agent's credence in the theory, which may be based on the evidence, without any taint of wishful thinking. Cognitive values function *as values*, not as surrogates for evidence.

2. Should Acceptance be Rejected?

There is a difficulty, however, with understanding, within a framework that represents belief states as probability assignments, what role there is for a notion of acceptance or rejection of hypotheses to play. This was the thrust of Jeffrey's response to Rudner (Jeffrey, 1956). He agrees with Rudner that, if it is the job of the scientist to accept or reject hypotheses, then the scientist *qua* scientist must make value judgments. He then argues that the sensitivity of criteria for acceptance or rejection to what is at stake makes an epistemic model that consists of modelling the epistemic state of a scientist merely in terms of propositions accepted or rejected too crude; one ought instead to treat of reasonable degrees of belief. On Jeffrey's approach, acceptance falls by the wayside.

If we do accept acceptance, should it be modelled as a change in belief state, say, a shift to a belief state that assigns credence 1 to the accepted propositions? If so, then this is a shift, in the absence of new evidence, away from the credence arrived at by considering all evidence. This, as Earman puts it in his discussion of the issue, "is nothing short of madness" (Earman, 1992, 194). If, however, acceptance is not a belief shift, this limits the role it will play in the practice of science. Consider two agents, Peter and Richard. Peter labels some hypotheses as accepted, whereas Richard employs no notion of acceptance. They have the same credences and the same utilities, except for any purely cognitive utility that Peter attaches to the act of acceptance. Since Peter and Richard have the same credences and utilities, they make the same decisions, insofar as practical utilities are concerned. Any impact that acceptance will have on decisions will stem from consideration of cognitive utility associated with the act of acceptance itself. But if a belief state is represented by a credence-function, then it is hard to see how Peter might be in an epistemically more valuable state than Richard. Acceptance appears to be a wheel turning idly.

At the level of abstraction on which the discussion between Jeffrey and Rudner took place, in which agents are idealized and cognitive limitations of agents are ignored, Jeffrey is entirely correct. There is, however, a role, albeit a limited one, for a notion of acceptance in a probabilistic account, when the picture is refined so as to take cognitive limitations into account. It is too much to ask of an agent of limited cognitive resources to carry around a full probability assignment in her head, for any rich set of propositions, and such a thing could not be readily communicated. There will be cases, however, in which, though an agent's credence in a proposition is less than unity, this difference will have negligible effect on decisions, and the agent, in any foreseeable situation, will act as if she accepts the proposition as definitely true. In such a case, she would do well to regard the proposition as accepted, without worrying too much about what her credence in it ought to be. It will also be simpler to report her attitude as one of acceptance, instead of trying to communicate a credence-level. Note that these are cases in which the proposition is so well supported that the sensitivity to what is at stake discussed by Rudner and Jeffrey is largely absent; it is only

in such cases that we wish to label some propositions accepted. This proposal is similar in spirit to that of Sargent (2009), who, developing a suggestion of Frankish (2004), introduces a notion of acceptance to relieve an agent of finite cognitive resources from the burden of doing expected utility calculations in cases when the cost of doing the calculation would exceed the benefit.

On any such approach to introducing acceptance into a context in which belief-states are represented by probability functions, acceptance will play a fairly limited role, and will be incapable of bearing the full burden that acceptance-based approaches to cognitive utility require of it. Instead of looking only to cognitive utility that attaches to the act of acceptance of a hypothesis, we should consider also *epistemic utilities*—utilities that attach to a state of belief. This is an approach that has been advocated by, among others, Oddie (1997), Joyce (1998), and Greaves and Wallace (2006). See Fallis (2007) for a superb overview.

3. Epistemic Utility

Though utility, cognitive or otherwise, has a role to play in any decision of whether to accept a hypothesis, this does not exhaust the role of utility in a scientific context. An investigation that leads to better-informed credences about a class of hypotheses may have value even if it does not lead to acceptance of any of them, if it results in a belief state that leads to better decisions or has some cognitive value above and beyond its impact on decisions. We must, therefore, be prepared to talk about the value of a belief state.

In this paper, an agent's state of belief will be represented by an assignment of numerical degrees of belief, or *credences*, to a set of propositions. It will be assumed that these credence-functions satisfy the axioms of probability. This is the basic starting point of probabilist approaches to issues of theory confirmation. No claim is made as to psychological accuracy of this picture. It should be noncontroversial that human beings do not possess precise numerical degrees of belief. Moreover, our qualitative rankings of propositions as more or less credible are often incompatible with the existence of numerical degrees of belief that represent these qualitative rankings and satisfy the axioms of probability. Nevertheless, it is possible to regard these facts about human beings as being attributable to our cognitive limitations. We may regard the model of an agent that has numerical credences as a *regulative ideal*: insofar as we depart from this model, this is to be regarded as a departure from ideal rationality. If one becomes aware that one's qualitative rankings of propositions as more or less credible are incompatible with the existence of quantitative degrees of belief satisfying the probability calculus, then one ought to regard such judgments as flawed and (if it matters) attempt to mend them.

Example 1. To motivate the idea of the value of a state of belief, consider first the following simple example. On Tuesday Bob is offered the option of purchasing tickets from a deck of 100 tickets, labeled 1 through 100. The ticket numbered n costs n dollars, and each ticket Bob purchases can be redeemed for a prize of 100 dollars on Wednesday, if it rains that day. We assume that Bob has the funds to purchase as many tickets as he deems fit, and that we can treat Bob's utilities as being linear in money.

Suppose, now, that Bob has degree of belief p that it will rain tomorrow. Then he will regard any ticket with price less than $\$100p$ as a good deal, and any ticket with price greater than $\$100p$ as a bad deal (It matters little whether he chooses to buy a ticket with price exactly equal to $\$100p$, should there be one; we will assume that he does.) Let n_p be the

least integer less than or equal to $100p$. In purchasing his tickets Bob expends a total of

$$\sum_{n=1}^{n_p} n = \frac{1}{2} n_p(n_p + 1).$$

In the event of rain Wednesday, Bob wins all his bets, for total winnings of $\$100n_p$, and so his net gain is

$$V_r(p) = 100 n_p - \frac{1}{2} n_p(n_p + 1).$$

If it doesn't rain, Bob has a net loss:

$$V_{\sim r}(p) = -\frac{1}{2} n_p(n_p + 1).$$

$V_r(p)$ is the value to Bob of having credence p in r , if r is true, and $V_{\sim r}(p)$ is the value of having that credence, if r is false.

Suppose, now, that Alice has degree of belief q in r . Then her assessment of the value to Bob of having credence p in r is her expectation value of the reward to Bob; that is, a weighted mean of $V_r(p)$ and $V_{\sim r}(p)$, where the weighting is her own credence:

$$U(p; q) = q V_r(p) + (1 - q)V_{\sim r}(p).$$

This is equal to $U(q; q)$ when $n_p = n_q$ (that is, when Bob makes exactly the same decisions as Alice), and is less than $U(q; q)$ if there are any tickets that Bob would buy and Alice would not, or *vice versa*. Therefore, we have, for any p, q ,

$$U(q; q) \geq U(p; q).$$

Example 2. Suppose that we want to rate weather forecasters on their forecasting ability. We devise the following scheme. Each forecaster is to state a probability p of rain the next day, and be awarded p points if it rains, $1 - p$ points if it doesn't. This rewards high values of p in the case of rain, low values in the case of no rain.

Though this may appear at first glance to be a good scoring rule, it is problematic, because it gives forecasters incentive to "work the system" by reporting probabilities other than their genuine degrees of belief. If Alice has degree of belief q that it will rain, then she judges the expected points awarded, if she reports p , as

$$q p + (1 - q) (1 - p) = (1 - q) + (2q - 1)p.$$

If $q = 1/2$, all choices of p have the same expected reward. But if $q > 1/2$, Alice maximizes her expected reward by choosing $p = 1$, and if $q < 1/2$, by choosing $p = 0$.

Example 3. A scoring rule that does not have this defect is the *Brier rule* (Brier, 1950). We reward $(1 - p)^2$ points if it rains, p^2 points if it doesn't rain, and count a low score as better than a high one. The expected point-score for a forecaster who reports p , as judged by an agent whose credence in rain is q , is

$$q(1 - p)^2 + (1 - q)p^2 = q(1 - q) + (p - q)^2,$$

which is obviously minimized by choosing $p = q$. Thus, forecasters will minimize their expected score by reporting their actual degrees of belief. Scoring rules that have this feature are called *proper scoring rules*.³

³ For some of the seminal work on proper scoring rules, see Good (1952), McCarthy (1956), Savage (1971).

Let \mathcal{S} be a set of mutually exclusive and jointly exhaustive hypotheses $\{h_i \mid i = 1, \dots, n\}$, and let $\Omega_{\mathcal{S}}$ be the set of probability functions on \mathcal{S} . An *epistemic utility function* \mathcal{V} is an assignment, to each $h_i \in \mathcal{S}$, of a function $V_i : \Omega_{\mathcal{S}} \rightarrow \mathbb{R}$, where $V_i(\mathbf{p})$ gives the value of having credences $\mathbf{p} = (p_1, p_2, \dots, p_n)$ if the hypothesis h_i is true. These might include, as in Example 1, concrete payoffs as a result of decisions made on the basis of those credences, or cognitive values, or both. The expected utility to Bob of having credence-function \mathbf{p} , as judged by Alice, who has credence-function \mathbf{q} , is

$$U^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) = \sum_{i=1}^n q_i V_i(\mathbf{p}).$$

Following Greaves and Wallace (2006), let us say that a credence-function \mathbf{q} is *self-recommending for* \mathcal{V} iff an agent with utilities \mathcal{V} and credence-function \mathbf{q} does not judge any other credence-function to be more valuable than her own, that is, iff

$$U^{\mathcal{V}}(\mathbf{q}; \mathbf{q}) \geq U^{\mathcal{V}}(\mathbf{p}; \mathbf{q})$$

for all $\mathbf{p} \in \Omega_{\mathcal{S}}$. \mathbf{q} is *strictly self-recommending for* \mathcal{V} iff an agent with epistemic utility function \mathcal{V} and credence-function \mathbf{q} judges every other credence-function to be strictly worse than her own, that is, iff

$$U^{\mathcal{V}}(\mathbf{q}; \mathbf{q}) > U^{\mathcal{V}}(\mathbf{p}; \mathbf{q})$$

for all $\mathbf{p} \neq \mathbf{q}$.⁴

An epistemic utility function \mathcal{V} is (*strictly*) *stable* at \mathbf{p} iff \mathbf{p} is (strictly) self-recommending for \mathcal{V} ; it is *everywhere (strictly) stable* iff it is (strictly) stable at every $\mathbf{p} \in \Omega_{\mathcal{S}}$. Everywhere stable epistemic utility functions correspond to proper scoring rules.

We define the quantity $\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q})$ as the difference, as assessed by \mathbf{q} , between the value of \mathbf{q} and the value of \mathbf{p} .

$$\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) = U^{\mathcal{V}}(\mathbf{q}; \mathbf{q}) - U^{\mathcal{V}}(\mathbf{p}; \mathbf{q}).$$

\mathcal{V} is stable at \mathbf{q} iff $\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) \geq 0$ for all \mathbf{p} , and strictly stable iff $\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) > 0$ for all $\mathbf{p} \neq \mathbf{q}$.

Example 4. *Generalized Brier rule.* Let $\{a_i\}$ be positive numbers, and define

$$V_i(\mathbf{p}) = -a_i(1 - p_i)^2 - \sum_{j \neq i} a_j p_j^2.$$

This gives

$$U^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) = - \sum_i a_i (q_i(1 - q_i) + (p_i - q_i)^2)$$

and

$$\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) = \sum_i a_i (p_i - q_i)^2.$$

Epistemic utility functions of this form are, therefore, everywhere strictly stable.

Example 5. One may want to consider epistemic utilities that depend only on the agent's degree of belief in the true hypothesis. It can be shown that, for $n > 2$, the only everywhere stable epistemic utility functions that have this feature are those of the form

$$V_i(\mathbf{p}) = a_i + \log_b(p_i),$$

⁴ Oddie (1997)'s term for a credence-function that is not self-recommending is "self-defeating" (p. 537).

where $\{a_i\}$ is an arbitrary set of numbers, and b is any number greater than 1.⁵ Note that, though V_i is a continuous function on the interior of Ω_S , it is unbounded as $p_i \rightarrow 0$, and so cannot be extended to a continuous function defined at $p_i = 0$. A demand for a stable epistemic utility function that is defined on the whole of Ω_S and has V_i depend only on p_i cannot be satisfied.

For this choice of epistemic utility function, with $b = 2$, we have

$$\Delta^V(\mathbf{p}; \mathbf{q}) = \sum_i q_i \log_2 (q_i/p_i),$$

which is the relative entropy of \mathbf{q} with respect to \mathbf{p} , also known as the Kullback-Leibler discrepancy, a commonly used measure of the difference between two probability functions.

It is easy to show that epistemic utility functions that evaluate credences only in terms of the consequences of decisions that might be made on the basis of them *must* be everywhere stable. Let \mathcal{A} be a set of acts, and let $U_i(A)$ be the utility of choosing act A if h_i . For a credence-function \mathbf{q} , let $A_{\mathbf{q}}$ be an act that maximizes expected utility. This gives us $V_i(\mathbf{q}) = U_i(A_{\mathbf{q}})$. To say that $A_{\mathbf{q}}$ maximizes expected utility as judged by \mathbf{q} means that

$$\sum_i q_i U_i(A_{\mathbf{q}}) \geq \sum_i q_i U_i(A)$$

for all $A \in \mathcal{A}$. The value of having credences \mathbf{p} , as judged by an agent whose credence-function is \mathbf{q} , is

$$U(\mathbf{p}; \mathbf{q}) = \sum_i q_i U_i(A_{\mathbf{p}}).$$

Because $A_{\mathbf{q}}$ maximizes expected utility, as judged by \mathbf{q} ,

$$\sum_i q_i U_i(A_{\mathbf{q}}) \geq \sum_i q_i U_i(A_{\mathbf{p}})$$

for all \mathbf{p} , and $U(\mathbf{q}; \mathbf{q}) \geq U(\mathbf{p}; \mathbf{q})$. This is true whether the utility that attaches to acts is cognitive or practical. Therefore, approaches to cognitive utility that locate all cognitive utility in acts of acceptance and rejection necessarily yield epistemic utility functions that are everywhere stable.

4. The Value of Learning Theorem

There is a theorem, first published by I. J. Good (1967), and generalized by Graves (1989), Skyrms (1990) (Ch. 4), and Oddie (1997), according to which learning can never have negative expected value. This is true even when there is a possibility of being misled, that is, when among the possible results of the learning experience is a move in the wrong direction, a move away from belief in the true hypothesis, perhaps resulting in worse decisions than those that would have been made on the basis of the pre-learning credences.⁶ The theorem

⁵ See Savage (1971), §9.4; McCarthy (1956) attributes this result to unpublished work of Andrew Gleason.

⁶ For instance, an unlucky series of tosses could lead someone to think a fair coin biased.

says that, even when this is possible, the *expected value* of the learning will not be negative; the possible negative consequences will be outweighed by the positive ones.

Good, Graves, and Skyrms evaluate learning in terms of its potential consequences for subsequent decisions. But the value of doing an experiment, or of pursuing some line of research, can be due, in part, to the expected increase in the value of one's belief state. For example, although some practical benefit might emerge from research performed at the Large Hadron Collider (LHC), the primary reason for building it is the value of the increased insight into the laws of nature to be gained from experimentation with the Collider. We don't know in advance what the outcome of experimentation with the LHC will be, of course. It may be that the Standard Model of particle physics is confirmed; it may be that wholly unexpected results are found, provoking a radical revision in fundamental theory. An estimate of the value of building the LHC should be based on some assessment (perhaps vague) of the *expected* cognitive utility of experimentation to be conducted with it.

Oddie (1997) showed that, for updating via conditionalization, the theorem holds if cognitive epistemic values are included, with the condition that the epistemic utility functions be everywhere strictly stable, a condition that Oddie calls *cogency*. The version we give here applies, as in Skyrms, to any learning experience, not just conditionalization, and will be phrased to highlight the minimal stability condition that is required. This will be important in the next section, in which this condition will be justified as a reasonable requirement on epistemic utility functions.

Suppose that Alice has an epistemic utility function \mathcal{V} , and that she is about to undergo an experience that may change her credences in the members of a partition $\{h_i\}$. Suppose that she will make a transition from her current credence-function cr to one of a set $\{cr_j\}$. This might be through Bayesian conditionalization on new information (as in the treatments of Good and Oddie), or through Jeffrey conditionalization (as in Graves), or some other process; we assume only, following Skyrms, that her current credence-function cr satisfies what Graves calls *condition M* (for Miller's Principle) and van Fraassen (1984) calls *Reflection*:⁷

$$cr(h_i \& e_j) = cr(e_j) cr_j(h_i),$$

where e_j is the proposition that the transition will be a transition to cr_j . Note: we are not taking it as a condition of rationality that *all* changes in credence satisfy condition *M*; rather, following Skyrms, we take condition *M* to be a necessary condition for the transition to count as a pure learning experience. Credence changes that don't satisfy condition *M* don't count as learning.

If $h_i \& e_j$, then the value of Alice's post-learning credence is $V_i(cr_j)$, and the epistemic value of having had the learning experience is

$$V_i(cr_j) - V_i(cr).$$

Alice therefore judges the expected value of undergoing the learning experience as

$$\sum_{i,j} cr(h_i \& e_j) (V_i(cr_j) - V_i(cr)).$$

Applying condition *M* this becomes

$$\sum_j cr(e_j) \sum_i cr_j(h_i) (V_i(cr_j) - V_i(cr))$$

⁷ In Skyrms' presentation, *M* is for Martingale (Skyrms, 1990, 105).

$$\begin{aligned}
&= \sum_j cr(e_j) (U^{\mathcal{V}}(cr_j; cr_j) - U^{\mathcal{V}}(cr; cr_j)) \\
&= \sum_j cr(e_j) \Delta^{\mathcal{V}}(cr; cr_j).
\end{aligned}$$

This is Alice's judgment, using her current credences, of the expected change in value of her credences due to the learning experience. However, because of condition M , it is also her expectation value of $\Delta^{\mathcal{V}}(cr; cr_j)$, which is the change in value of her credence, as judged by the new credence.

If \mathcal{V} is stable at cr_j , then $\Delta^{\mathcal{V}}(cr; cr_j) \geq 0$. Therefore, if, for every j such that $cr(e_j) \neq 0$, \mathcal{V} is stable at cr_j , each term of the sum is non-negative, and Alice's evaluation of the expected value of undergoing the learning experience is non-negative. It is strictly positive if, in addition, $\Delta^{\mathcal{V}}(cr; cr_j) > 0$ for some j such that Alice has non-zero credence that she will end up with credence cr_j . Since we have shown in the previous section that practical epistemic utility functions are everywhere stable, we get Skyrms' version of the theorem as a special case.

Of course, in real-life learning situations there will typically be costs associated with undergoing the learning experience, if only an opportunity cost due to the time spent learning. The Value of Learning Theorem does not show that learning will always have non-negative utility, all things considered. It concerns only the contribution to the expected utility calculation due to the change in credence.

Note that it is condition M that turns Alice's current expectation of the change in value of her belief state into an expectation of the change in value, as judged by the possible new belief states. This is why a parallel argument would not establish that it would be rational for an agent, using her current credences to make the assessment, to take a drug to change her belief state, even though the new belief state would be judged an improvement by the new belief state. Thus, Fallis misses the mark when he says,

In order to decide whether it is rational to take this drug, we have to use \mathbf{r}_d [the drug-induced credence function] to calculate the expected epistemic utility of her doxastic state if she takes this drug.... 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 [epistemic utility functions] 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 (Fallis, 2007, 231).

The requirement that belief states not be changed absent any new information is not needed as an extra requirement, if stability is already required. This is because the decision of whether to take the drug is evaluated on the basis of the agent's credences at the time the decision is made. In situations like the drug-taking decision, in which condition M will not be satisfied, the expected utility of making the credence-shift will be negative.

The condition that \mathcal{V} be stable at the credences that could result from the learning process, though sufficient, is not a necessary condition for a learning process to have positive expected value. Horwich (1982, 127–28) utilized, for a two-element partition, the utility function of Example 2, above. Using this function, he proved that learning has positive expected utility when the evidence to be acquired is not irrelevant to the hypothesis in question.

5. Are all reasonable epistemic utilities stable?

As we have seen, practical epistemic utility functions—that is, epistemic utility functions that evaluate a belief-state in terms of its effects on subsequent decisions—are necessarily everywhere stable. We have not thus far required this of cognitive epistemic utility functions.

Patrick Maher has pointed out that the Value of Learning Theorem does not hold for arbitrary epistemic utility functions, or even for arbitrary truth-valuing functions, where, for a two-element partition, $\{h, \sim h\}$, an epistemic utility function \mathcal{V} is *truth-valuing*⁸ iff V_h increases, and $V_{\sim h}$ decreases, as degree of belief in h increases (this concept is generalized in §6, below).

That an epistemic utility function be truth-valuing is not a sufficient condition for the Value of Learning Theorem to hold for it—or, as Maher calls it, the Scientific Value of Evidence Thesis (SVET). This fact, which Maher illustrates with a simple example (Maher, 1993, 176–77), he uses to argue for the inferiority of probabilist explanations of SVET—that is, explanations in terms of value that accrues to belief-states, represented as probability assignments—to acceptance-based explanations. He considers the possibility of placing restrictions on epistemic utility functions, but rejects such a move as *ad hoc*.

Another possible response is to say that a utility function simply does not count as representing the cognitive goals of science unless it makes SVET true. ... The proposed restriction on what counts as a scientific utility function is not motivated by any reflection on the goals of science that the utility function is supposed to capture, but merely by the fact that the probabilist explanation fails without this restriction; thus the probabilist restriction on what counts as a scientific utility function is quite *ad hoc* (Maher, 1993, 177–78).

We should ask whether one can obtain the SVET without any moves that are objectionably *ad hoc*.

Let us consider an agent who has the choice of obtaining new evidence, or not, and who has some epistemic utility function. Without any assumptions about the form this epistemic utility function takes, we cannot conclude that, for her, conditionalizing on the new evidence will have non-negative expected epistemic utility. Maher asks: why should we expect her epistemic utilities be such as ensure that conditionalizing *will* have non-negative expected epistemic utility?

One way to approach the question is to assume that her belief-shift upon learning new evidence will be the shift brought about by conditionalizing on the new evidence, and ask whether we can justify a constraint on her epistemic utility function that guarantees that conditionalizing has non-negative expected epistemic utility. Suppose, however, that our agent's epistemic utilities do not satisfy any such constraint, and that, for her, conditionalization has negative expected utility. For such an agent, it is unreasonable to assume that she will conditionalize; this belief-updating strategy will be regarded by her as inferior, in the sense of having lower expected epistemic utility, to the strategy of simply leaving her credences unchanged upon acquiring new evidence. It seems more reasonable to assume that she will adopt some updating strategy that maximizes expected epistemic utility.

This may look as if it trivializes the Scientific Value of Evidence Thesis; if we assume that our agent adopts an updating strategy that maximizes expected epistemic utility, the expected epistemic utility of acquiring new evidence is guaranteed to be non-negative and will be strictly positive except in cases in which *no* updating strategy has positive expected

⁸ Maher's term is "truth-seeking."

epistemic utility. We can still ask, however: why should *conditionalization* have positive expected epistemic utility?

Greaves and Wallace (2006) have shown that conditionalization maximizes expected utility if and only if, for every e_j such that $cr(e_j) > 0$, \mathcal{V} is stable at the conditional credence-function $cr(\cdot | e_j)$. Requiring that epistemic utility functions be stable, therefore, suffices to ensure that our utility-maximizer will conditionalize.

Oddie (1997) argues for stability on the basis of *conservatism*, the principle that “[a]bsent any new information you should not change your cognitive state” (p. 538). Oddie is not alone in finding this a compelling principle; we have already quoted Earman to the effect that violation of this principle is “nothing short of madness”. For those who agree with Oddie and Earman, this ought to be sufficient; we require that epistemic utility functions be stable and thereby obtain the Value of Learning result. Maher, on the other hand, regards stipulation of stability as another *ad hoc* move. After pointing out the instability of the epistemic utility function considered by Horwich (see our Example 2, above), Maher notes that a possible response is “to say that a utility function does not count as representing the cognitive goals of science unless it is a proper scoring rule.” However, “like the similar move used to make SVET necessary, success is here bought at the price of invoking a completely *ad hoc* assumption” (Maher, 1993, 179).

If, however, we have independent justification for conditionalization, then this will become an argument in favour of a stability condition; if conditionalization does not maximize expected epistemic utility, our epistemic utility maximizer will run afoul of the reasons we have for thinking that she ought to conditionalize. One such justification is the diachronic dutch book argument. De Finetti and Ramsey showed that $c(h|e) = cr(h\&e)/cr(e)$ is the fair price of a bet on h , conditional on e , where a bet conditional on e is one that is called off, and all monies returned, if $\sim e$. The diachronic dutch book argument (reported in Teller (1973), and attributed therein to David Lewis) uses this to show that an agent who updates via some strategy other than conditionalization is vulnerable to a diachronic dutch book. This means that there is a set of bets, some offered prior to updating, some after, that the agent will regard as all favourable, whose net effect is a sure loss.

Suppose our partition $\{h_i\}$ consists of propositions that are wagerable—that is, propositions whose truth-value can be ascertained, with sufficient reliability, in a short enough time, that it makes sense to place bets on them. An agent whose belief-shift maximizes expected epistemic utility will leave herself open to a diachronic dutch book, unless her epistemic utility function is stable at each credence-function obtainable from her credence by conditionalizing on some $\{e_j\}$ with $cr(e_j) > 0$. Avoidance of a diachronic dutch book is a reasonable constraint to place on a rational agent undergoing learning. Thus, for cases like this, in which the hypotheses under question are wagerable, an epistemic utility function that can represent the considered value judgments of a reasonable agent must be stable at every point in probability-space that could be reached from her current credences by conditionalization on evidence that she deems it possible to acquire. This condition is sufficient for the Value of Learning Theorem of the previous section to apply.

Not all cases of interest will be of this sort, however. Scientific theories, if they involve putative laws of nature, make claims with spatiotemporally unbounded import, and hence cannot be conclusively verified by spatiotemporally bounded observation. For this reason, there may be no satisfactory way to set payoff conditions, and no sensible way of making bets on the hypotheses.

Even if the hypotheses in question are not wagerable, precluding a diachronic dutch book, there is still an issue of diachronic consistency raised by unstable epistemic utilities. Suppose

that an agent has an epistemic utility function that is not stable at $cr(\cdot|e_j)$. Then a strategy that maximizes expected epistemic utility will recommend some other credence function cr_j , not identical to $cr(\cdot|e_j)$, upon learning e_j . This means that, though her credence that she will make a transition to cr_j is $cr(e_j)$, there is some proposition h such that

$$cr_j(h) \neq cr(h \& e_j)/cr(e_j).$$

That is, her credences violate condition M .

Condition M cannot be plausibly be taken as a constraint on all rational belief-changes. But it *can* be taken as a condition for a belief-change to count as a learning experience, one in which the post-shift agent endorses her prior judgments as reasonable for someone in her prior epistemic condition. An agent who maximizes expected epistemic utility in her belief-shifts cannot satisfy condition M —that is, cannot undergo a pure learning experience—unless her epistemic utility function is suitably stable. An epistemic utility function that can represent the considered value-judgments of a reasonable agent ought not to foreclose the possibility of learning. This means that the epistemic utility function ought to be stable at every credence that can be reached from the agent’s current credence-function by conditionalization on evidence regarded by the agent as possible for her to acquire. If this set of credence functions reaches the whole of probability space, then we have an argument for stability everywhere.⁹ If not, we have argued for a condition that is weaker than stability everywhere. But it is still sufficient for the Value of Learning Theorem. Imposing this stability requirement, a reasonable one, and not at all *ad hoc*, we get the result that obtaining evidence has non-negative utility, and has positive utility if \mathcal{V} is strictly stable at some credence-function that is different from the agent’s current one and might be reached by conditionalizing on the evidence. We obtain the Scientific Value of Evidence Thesis for epistemic utilities that include cognitive utilities, without any objectionable moves.

Some might object that it is nonsensical to place constraints on values: an agent values whatever she values, and considerations of rationality or reasonableness cannot be applied in a critique of values, but enter only when consideration is made of how the agent is to best reach the goals prescribed by her values. To this, it should be pointed out that decision theory begins with constraints on preferences between acts, motivated by rationality concerns, and this does not seem nonsensical.

Moreover, advocates of an acceptance-based notion of cognitive utility also place constraints on these utilities. Maher proposes four necessary conditions for a cognitive utility function to be counted as reflecting scientific values (Maher, 1993, Ch. 9). The first of these is respect for truth, the condition that the cognitive utility of accepting a hypothesis when it is false not be higher than the utility of accepting it when it is true. This does, indeed, seem to be a reasonable condition to place on a cognitive utility function. The analogue, in the present context, would be to require that an epistemic utility function be truth-valuing. This, interestingly enough, turns out to be a consequence of stability.

⁹ There is a minor wrinkle, in that the set of experiences that can be distinguished by the agent may only be countable. In such a case, if the set of credence-functions reachable by conditionalization on possible evidence is dense in Ω_S , we can get everywhere stability by imposing a continuity condition on the utility function—provided that it extends to a continuous function on Ω_S .

6. Externalism and Truth-Valuing

Oddie (1997) distinguishes between *internalist* and *externalist* epistemic value functions. An internalist epistemic value function evaluates a belief state in terms of that state's structure, independently of how the world may be. It might, for example, value certainty over uncertainty, and rank credence-functions according to some measure of how concentrated they are. In our notation, internalism is the condition that $V_i(\mathbf{p})$ be the same function for each i . It is easy to see that, if this is the case, every credence function \mathbf{q} will assign the same value $U^V(\mathbf{p}; \mathbf{q})$ to \mathbf{p} , and, except for the trivial function that assigns the same value to every \mathbf{p} , no internalist value function can be everywhere stable, and no internalist value function whatsoever can be everywhere *strictly* stable (Oddie, 1997, 537).

This shows that an everywhere strictly stable epistemic value function must be an externalist one; it must be sensitive to the way the world is. Sensitive, but in what way? We would expect that an epistemic value function would value having high credence in true propositions; must this be imposed as an extra condition?

Fallis (2007) has shown that an everywhere stable epistemic utility function must satisfy a condition that he calls *weak monotonicity*, which is a form of truth-valuing. To understand what has been proven, we first need to discuss what is to be meant by a truth-valuing epistemic utility function.

For a partition consisting of only two hypotheses $\{h_1, h_2\}$, there's no difficulty in saying when an epistemic utility function is truth-valuing. In such a case $\Omega_{\mathcal{S}}$ is one-dimensional, so any change in credence is a change away from one vertex of $\Omega_{\mathcal{S}}$ and towards the other. We say that V_i is truth-valuing if $V_i(\mathbf{p})$ does not decrease as p_i is increased, and strictly truth-valuing if it increases as p_i is increased.

For $n > 2$, things are a bit more complicated. Let $\{\mathbf{e}_i\}$ be the vertices of the simplex $\Omega_{\mathcal{S}}$, with \mathbf{e}_i being the probability function that assigns probability 1 to h_i and 0 to all other members of \mathcal{S} . We will want to say that V_i is strictly truth-valuing if $V_i(\mathbf{p})$ increases as \mathbf{p} moves towards the vertex \mathbf{e}_i . And that means that we'll have to be a bit more precise about what it means to move towards the vertex, since, in a space of more than one dimension, there are different paths that can be taken from one point to another.

There are a number of metrics, or distance-functions, that can be defined on $\Omega_{\mathcal{S}}$. One is the *norm distance*:

$$d(\mathbf{p}, \mathbf{q}) = \frac{1}{2} \sum_{i=1}^n |p_i - q_i|.$$

This is the maximum difference in probabilities assigned by \mathbf{p} and \mathbf{q} to any of the 2^n propositions that can be formed from Boolean combinations of the elements of \mathcal{S} , and hence a natural measure of the difference between two probability assignments. The norm distance of any probability function \mathbf{p} from the vertex \mathbf{e}_i is $1 - p_i$.

Another metric on $\Omega_{\mathcal{S}}$ is, of course, the familiar Euclidean metric,

$$r(\mathbf{p}, \mathbf{q}) = \sqrt{\sum_i (p_i - q_i)^2}.$$

Generalizing, we define, for any set $\mathcal{W} = \{a_i\}$ of positive numbers, a weighted Euclidean metric,

$$r_{\mathcal{W}}(\mathbf{p}, \mathbf{q}) = \sqrt{\sum_i a_i (p_i - q_i)^2}.$$

To any metric on $\Omega_{\mathcal{S}}$ defining a notion of distance between two probability functions, there corresponds a notion of truth-valuing: V_i is truth-valuing at a point \mathbf{p} iff there is a neighbourhood of \mathbf{p} in which V_i does not decrease as the distance to \mathbf{e}_i decreases, and it is strictly truth-valuing if it increases with decreasing distance to \mathbf{e}_i . V_i is everywhere (strictly) truth-valuing iff it is (strictly) truth-valuing at each point in $\Omega_{\mathcal{S}}$, and the epistemic utility function \mathcal{V} is (strictly) truth-valuing iff every V_i is.

The notions of truth-valuing generated by different metrics need not coincide. The generalized Brier rule evaluates a credence function according to the square of a weighted Euclidean distance from the truth, and hence is strictly truth-valuing in that metric. It is not, however, truth-valuing in the norm distance (see Fallis 2007, 234).

We can also define a notion of truth-valuing that is independent of any consideration of a metric on $\Omega_{\mathcal{S}}$. This makes use of the fact that there is a natural affine structure on a space of probability functions, independent of any notion of distance. For distinct points \mathbf{p} , \mathbf{q} , we define the line containing \mathbf{p} and \mathbf{q} as the set of probability functions

$$(1 - t) \mathbf{p} + t \mathbf{q},$$

where t ranges over all values compatible with remaining in $\Omega_{\mathcal{S}}$. Call this line $\lambda(\mathbf{p}, \mathbf{q})$. We will require of any metric on $\Omega_{\mathcal{S}}$ that it be compatible with the affine structure; that is, for any admissible metric, it must be the case that the path along $\lambda(\mathbf{p}, \mathbf{q})$ is a path of minimal length among paths that join \mathbf{p} and \mathbf{q} .

For any point $\mathbf{p} \in \Omega_{\mathcal{S}}$, and any vertex \mathbf{e}_i , the direct path towards \mathbf{e}_i is along the line $\lambda(\mathbf{p}, \mathbf{e}_i)$. We will say that V_i is *affinely truth-valuing* at \mathbf{p} iff there is a neighbourhood of \mathbf{p} in which V_i does not decrease as we move towards the vertex \mathbf{e}_i along $\lambda(\mathbf{p}, \mathbf{e}_i)$, strictly affinely truth-valuing if it increases along this line. This is the condition that Fallis calls the *weak monotonicity constraint* (the condition that an epistemic utility function be everywhere truth-valuing in norm distance he calls the *monotonicity constraint*).

Since, on any admissible metric, the distance between \mathbf{p} and \mathbf{e}_i will decrease as we move towards \mathbf{e}_i along $\lambda(\mathbf{p}, \mathbf{e}_i)$, affine truth-valuing is weaker than any of the notions obtained from distance metrics. To be affinely truth-valuing is the weakest possible sense of truth-valuing; if an epistemic utility function is not affinely truth-valuing, it is clear that we should say that is not truth-valuing. The condition that an epistemic utility function be affinely truth-valuing captures the idea that the utility function be truth-valuing *in some sense*. The various metric-induced conditions impose further restrictions on what is to count on truth-valuing. In particular, truth-valuing in norm-distance entails that, if V_i is a continuous function, it depends only on p_i ; that is, the value of a belief state depends only on its degree of belief in the true hypothesis, and is independent of the way in which credences are distributed among false hypotheses.

There can be epistemic utility functions that are stable at some point, without being truth-valuing. Take, for example,

$$V_i(\mathbf{p}) = -p_i^2.$$

This epistemic utility function is nowhere truth-valuing, in any sense, as $V_i(\mathbf{p})$ decreases with increasing p_i . It is, however, strictly stable at the probability-function that assigns equal probability $1/n$ to each element of \mathcal{S} .

Fallis (2007) proves the lovely result that any epistemic utility function that is everywhere stable is affinely truth-valuing (this is his Proposition 1). Though Fallis does not express it this way, his result stems from the way $\Delta^{\mathcal{V}}$ behaves under mixtures. Take two probability functions \mathbf{p} , \mathbf{q} , and, for some $\alpha \in [0, 1]$, let

$$\mathbf{r} = \alpha \mathbf{p} + (1 - \alpha) \mathbf{q}.$$

Then, for any epistemic value function \mathcal{V} ,

$$\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r}) + \alpha \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p}) = (1 - \alpha) (\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) - \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{q})). \quad (*)$$

This relation can be readily verified by the reader.

Now take two points \mathbf{p} , \mathbf{r} , on a line that includes a vertex \mathbf{e}_i , with $p_i < r_i < 1$. Then \mathbf{r} is a mixture of \mathbf{p} and \mathbf{e}_i , and so

$$\mathbf{r} = \alpha \mathbf{p} + (1 - \alpha) \mathbf{e}_i$$

for some $\alpha \in (0, 1)$. This gives us $(1 - r_i) = \alpha(1 - p_i)$ and $(r_i - p_i) = (1 - \alpha)(1 - p_i)$. Putting \mathbf{e}_i for \mathbf{q} in (*), we get

$$\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r}) + \alpha \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p}) = (1 - \alpha) (\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{e}_i) - \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{e}_i)).$$

But

$$\begin{aligned} \Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{e}_i) - \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{e}_i) &= V_i(\mathbf{e}_i) - V_i(\mathbf{p}) - (V_i(\mathbf{e}_i) - V_i(\mathbf{r})) \\ &= V_i(\mathbf{r}) - V_i(\mathbf{p}), \end{aligned}$$

and so we have

$$\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r}) + \alpha \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p}) = (1 - \alpha) (V_i(\mathbf{r}) - V_i(\mathbf{p})).$$

Using $(1 - r_i) = \alpha(1 - p_i)$ and $(r_i - p_i) = (1 - \alpha)(1 - p_i)$, we get

$$(1 - p_i) \Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r}) + (1 - r_i) \Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p}) = (r_i - p_i) (V_i(\mathbf{r}) - V_i(\mathbf{p})). \quad (\dagger)$$

Note that this is unchanged when the roles of \mathbf{p} and \mathbf{r} are switched; therefore, the relation (\dagger) holds for *any* \mathbf{p} , \mathbf{r} lying on a line that extends to the vertex \mathbf{e}_i .

From (\dagger) the connection between stability and being affinely truth-valuing is immediate. If \mathcal{V} is stable at both \mathbf{r} and \mathbf{p} , both $\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r})$ and $\Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p})$ are non-negative, and so it follows that

$$(r_i - p_i) (V_i(\mathbf{r}) - V_i(\mathbf{p})) \geq 0.$$

If $r_i > p_i$, therefore, we must have $V_i(\mathbf{r}) \geq V_i(\mathbf{p})$. If \mathcal{V} is strictly stable at \mathbf{p} or \mathbf{r} , then $\Delta^{\mathcal{V}}(\mathbf{p}; \mathbf{r})$ or $\Delta^{\mathcal{V}}(\mathbf{r}; \mathbf{p})$ is strictly positive, and

$$(r_i - p_i) (V_i(\mathbf{r}) - V_i(\mathbf{p})) > 0.$$

Therefore, if $r_i > p_i$, $V_i(\mathbf{r}) > V_i(\mathbf{p})$. We conclude that an epistemic utility function is everywhere affinely truth-valuing if it is everywhere stable, and that it is everywhere strictly affinely truth-valuing if it is everywhere strictly stable. So, not only does stability (Oddie's 'cogency') entail sensitivity to the external word, it entails valuing truth.

7. Are there further restrictions on Epistemic Utility Functions?

Oddie (1997) imposes the condition that an epistemic utility function be everywhere strictly stable, on the grounds that it is always irrational to change one's credence in the absence of new evidence. This is a stipulation that Maher ((1990; 1993)) rejected as *ad hoc*.

In some discussions of theory choice, there seems to be a suggestion that considerations of cognitive value have a legitimate role to play in setting credences.¹⁰ This is only possible for epistemic utility functions that are not everywhere stable. An agent with a stable epistemic utility function will not be moved to adjust her credences by considerations of cognitive value.¹¹ Rather than ruling out by fiat the possibility of permitting considerations of cognitive value to affect our credences, we should ask whether an argument can be given for excluding this possibility. The considerations of §5 show that unstable epistemic value functions can run into problems of diachronic coherence, but they fall short of an argument that an epistemic value function must be everywhere stable. The argument only extends to stability at every point in probability space that can be reached by conditionalization on possible evidence. One could imagine, for example, that among the hypotheses in \mathcal{S} are two hypotheses $\{h_1, h_2\}$ such that the likelihoods $cr(e|h_1)$ and $cr(e|h_2)$ are the same for all evidence e that the agent deems it possible to obtain. This means that conditionalizing on evidence will not change the ratio of $cr(h_1)$ to $cr(h_2)$ (this is the probabilistic version of underdetermination). This is the sort of case in which epistemic utility might be called upon to set this ratio. One could imagine an epistemic utility function that is stable only for credence functions that assigned a certain value α to $cr(h_1)/cr(h_2)$.¹² An agent whose credences did not have $cr(h_1)/cr(h_2) = \alpha$ would be moved to change them to one that did. But, once achieved, this condition would be preserved under conditionalization on possible evidence, and so the agent would not run afoul of the considerations of §5. It is worthwhile

¹⁰ This statement is hedged because, though this is often *suggested* by such discussions (including Kuhn's), one rarely finds a clear statement, for the simple reason that it is often not clear *what* choices are under consideration. Putnam, for example, writes that “a great number of theories *must* be rejected on non-observational grounds, for the rule ‘Test every theory that occurs to anyone’ is impossible to follow” (Putnam, 2002, 140). Judgments of theoretical virtues such as coherence and simplicity are, according to Putnam, the grounds on which some theories are rejected out of hand. “In short, judgments of coherence, simplicity, and so on are presupposed by physical science. Yet coherence, simplicity, and the like are values” (Putnam, 2002, 142). It is not clear from this whether rejection entails low credence, or whether we are to choose not to bother with theories that lack the requisite virtues, whatever our credence in them might be.

¹¹ It might be suggested that there is a role for epistemic utilities in setting priors, even if the epistemic utility functions are everywhere stable. Instead of imagining an agent who already has credences and is considering whether it is in her interest to change them, we might consider the problem of endowing a *tabula rasa* with credences in the first place. On this sort of proposal, prior credences might be ranked according to $U(\mathbf{q}; \mathbf{q})$. Unfortunately, for everywhere strictly stable epistemic utility functions, picking a prior \mathbf{q} that maximizes $U(\mathbf{q}; \mathbf{q})$ leads to dogmatism. Since every V_i achieves its maximum at \mathbf{e}_i , the maximum value of $U(\mathbf{q}; \mathbf{q})$ is achieved at the vertex that maximizes $V_i(\mathbf{e}_i)$.

¹² Are there any such functions? It is easy to show that such functions exist, though whether there are any that are plausible is another matter. Let \mathcal{V}^S be any everywhere strictly stable epistemic utility function, and let \mathcal{T} be an invertible transformation of $\Omega_{\mathcal{S}}$ that leaves only probability functions with the desired condition invariant. Define \mathcal{V} by

$$V_i(\mathbf{p}) = V_i^S(\mathcal{T}^{-1}\mathbf{p}).$$

Then $U^{\mathcal{V}}(\mathbf{p}; \mathbf{q}) = \sum_i q_i V_i^S(\mathcal{T}^{-1}\mathbf{p})$ is uniquely maximized by $\mathbf{p} = \mathcal{T}\mathbf{q}$, and so \mathcal{V} is stable only on the subset of $\Omega_{\mathcal{S}}$ of invariant points of \mathcal{T} .

to ask whether some other argument could be given for everywhere stability, that did not simply presume that it is always wrong to let considerations of epistemic utility affect one's credences. An open question: is there a No Wishful Thinking theorem?

Besides stability, another constraint that might be considered is that epistemic utility functions be truth-valuing according to one's favourite metric. Fallis considers imposing the condition that epistemic utility functions be truth-valuing in norm distance, on the grounds that this captures the requirement that scientist should "seek the truth and nothing but the truth" (Fallis, 2007, §7). This is a strong constraint. In particular, it entails that $V_i(\mathbf{p})$, if it is a continuous function, depends only on p_i . The only epistemic utility functions that satisfy this condition and are strictly stable are of the form,

$$V_i(\mathbf{p}) = a_i + \log_b p_i,$$

where $b > 1$. This, according to Fallis, is an unreasonably tight restriction.

The logarithmic rule is almost certainly a scientifically acceptable EUF. In fact, as I discuss below, it has several very attractive properties. 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. 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 (Fallis, 2007, 237).

Fallis concludes from this that "Bayesian models of epistemic utilities fail as normative accounts of scientific inquiry" (240).

This is a strong conclusion. Is it warranted? It is certainly true that the requirement that $V_i(\mathbf{p})$ depend only on p_i is more constraining than might at first appear. One possible reaction to this would be to conclude that a seemingly reasonable requirement that leads to such consequences is not as innocent as it appeared, and that it is in fact too restrictive a condition. This seems to be the right reaction; there are no compelling arguments for the requirement, and respect for the truth does not require one to be indifferent as to how one's credences are distributed among false hypotheses. Another reaction might be to "bite the bullet," as Fallis says (237), and accept that the logarithmic rule is the only scientifically acceptable epistemic utility function. If need be, the Bayesian's teeth could be strengthened by arguments for the logarithmic rule that rest on considerations other than stability (see, *e.g.*, Good 1973, 116–118). This would be the appropriate reaction if there were good reasons to require that $V_i(\mathbf{p})$ be truth-valuing in norm distance.

An analogy: imagine someone who found it *prima facie* plausible to suppose that there are everywhere strictly stable epistemic utility functions that are not truth-valuing, on the grounds that it is not clear why a scientist who is content with her own epistemic state should be required to value belief in the truth. This person is then confronted with Fallis' theorem, and is surprised to learn that no such utility functions exist. The reasonable response, it would seem, would be to abandon the thought that there should be such utility functions, and that a requirement that there be such functions, despite its *prima facie* plausibility, is in fact impossibly strong. It would be unreasonable to throw out the framework.

8. Conclusion

In §5, it was argued that epistemic utility functions that can represent the considered judgments of a reasonable agent ought to be stable at each point in probability-space that could be reached from the agent's current credences by conditionalization on possible evidence. It

would be good to have an argument for a stronger claim, that there is something unreasonable about having epistemic values that are not everywhere stable, but such an argument, it seems, would have to proceed along entirely different lines. Nevertheless, the stability condition argued for is strong enough to establish the Value of Learning Theorem, opening the way for consideration of cognitive epistemic values, in addition to values associated with acceptance and rejection of hypotheses, to play a motivating factor in science.

9. Acknowledgments

I would like to thank John Nicholas, Don Fallis, and an anonymous referee for helpful comments. This work was supported by a grant from the Social Sciences and Humanities Research Council of Canada.

References

- Brier, G. W.: 1950, 'Verification of Forecasts Expressed in Terms of Probability'. *Monthly Weather Review* **78**, 1–3.
- Douglas, H.: 2000, 'Inductive Risk and Values in Science'. *Philosophy of Science* **67**, 559–579.
- Earman, J.: 1992, *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: The MIT Press.
- Fallis, D.: 2007, 'Attitudes Toward Epistemic Risk and the Value of Experiments'. *Studia Logica* **86**, 215–246.
- Frankish, K.: 2004, 'Partial Belief and Flat-Out Belief'. In: F. Huber and C. Schmidt-Petri (eds.): *Degrees of Belief: An Anthology*. Oxford: Oxford University Press, pp. 75–93.
- Good, I. J.: 1952, 'Rational Decisions'. *Journal of the Royal Statistical Society, Series B* **14**, 107–14.
- Good, I. J.: 1967, 'On the Principle of Total Evidence'. *The British Journal for the Philosophy of Science* **17**, 319–322.
- Good, I. J.: 1973, 'Information, Rewards, and Quasi-Utilities'. In: J. Leach, R. Butts, and G. Pearce (eds.): *Science, Decision, and Value*, Western Ontario Series in the Philosophy of Science, Vol. 1. D. Reidel Publishing Company, pp. 115–127.
- Graves, P. R.: 1989, 'The Total Evidence Theorem for Probability Kinematics'. *Philosophy of Science* **56**, 317–324.
- Greaves, H. and D. Wallace: 2006, 'Justifying conditionalization: Conditionalization maximizes expected epistemic utility'. *Mind* **115**, 607–32.
- Hempel, C. G.: 1965, 'Science and Human Values'. In: *Aspects of Scientific Explanation*. New York: Macmillan Publishing Co., pp. 81–96.
- Hempel, C. G.: 1981, 'Turns in the Evolution of the Problem of Induction'. *Synthese* **46**, 389–404. Reprinted in J. H. Fetzer, ed., *The Philosophy of Carl G. Hempel*, Oxford University Press, 2001, 344–371.
- Horwich, P.: 1982, *Probability and Evidence*. Cambridge University Press.
- Jeffrey, R.: 1956, 'Valuation and Acceptance of Scientific Hypotheses'. *Philosophy of Science* **23**, 237–46. Reprinted in *Probability and the Art of Judgment*, Cambridge University Press, 2002, 14–29.
- Joyce, J.: 1998, 'A Nonpragmatic Vindication of Probabilism'. *Philosophy of Science* **65**, 575–603.
- Kuhn, T. S.: 1977, 'Objectivity, Value Judgment, and Theory Choice'. In: *The Essential Tension*. Chicago: The University of Chicago Press, pp. 320–339.
- Levi, I.: 1967, *Gambling With Truth: An Essay on Induction and the Aims of Science*. Cambridge, MA: The MIT Press.
- Longino, H.: 1995, 'Gender, Politics, and the Theoretical Virtues'. *Synthese* **104**, 383–397.
- Maher, P.: 1990, 'Why Scientists Gather Evidence'. *The British Journal for the Philosophy of Science* **41**, 103–119.
- Maher, P.: 1993, *Betting on Theories*. Cambridge University Press.
- McCarthy, J.: 1956, 'Measures of the Value of Information'. *Proceedings of the National Academy of Sciences of the United States of America* **42**, 654–655.
- Oddie, G.: 1997, 'Conditionalization, Cogency, and Cognitive Value'. *The British Journal for the Philosophy of Science* **48**, 533–541.
- Putnam, H.: 2002, *The Collapse of the Fact/Value Dichotomy and Other Essays*. Harvard University Press.

- Rudner, R.: 1953, 'The Scientist *Qua* Scientist Makes Value Judgments'. *Philosophy of Science* **20**, 1–6.
- Sargent, M.: 2009, 'Answering the Bayesian Challenge'. *Erkenntnis* **70**, 237–252.
- Savage, L. J.: 1971, 'Elicitation of Personal Probabilities and Expectations'. *Journal of the American Statistical Association* **66**, 783–801. Reprinted in Henry E. Kyburg and Howard E. Smokler, eds., *Studies in Subjective Probability* (Robert E. Krieger Publishing Company, 1980), 149–191.
- Skyrms, B.: 1990, *The Dynamics of Rational Deliberation*. Cambridge, MA: Harvard University Press.
- Teller, P.: 1973, 'Conditionalization and Observation'. *Synthese* **26**, 218–258.
- van Fraassen, B.: 1984, 'Belief and the Will'. *Journal of Philosophy* **81**, 235–256.