

Ridgeview Publishing Company

The Value of Truth

Author(s): Barry Loewer

Source: *Philosophical Issues*, Vol. 4, Naturalism and Normativity (1993), pp. 265-280

Published by: Ridgeview Publishing Company

Stable URL: <https://www.jstor.org/stable/1522845>

Accessed: 12-06-2019 18:23 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Ridgeview Publishing Company is collaborating with JSTOR to digitize, preserve and extend access to *Philosophical Issues*

The Value of Truth

Barry Loewer

It is a philosophical commonplace that truth is the primary norm for belief and assertion. For example, Frege held that truth is the aim of science and Dummett that “it is part of the concept of truth that we aim at making true statements”.¹ Contrary to this commonplace Steve Stich has recently argued that “once we have a clear view of the matter, most of us will not find *any value*, either intrinsic or instrumental, in having true beliefs”.² This is really an astonishing claim. It seems to entail that it is not rational to seek or take steps to insure that our beliefs are true. And if that is right then the effort we expend with the purpose of discovering the truth are expended under an illusion. Further, widely held views in epistemology (e.g. that knowledge = justified true belief is worth having and that our beliefs ought track the truth), the philosophy of logic (that reasoning ought to conform to truth preserving rules), the philosophy of science (that scientists ought attempt to find true theories) and ethics (that lying is morally wrong) are put in doubt. If Stich is right in claiming that we will (ought) not find *any value* in true belief then presumably we ought to alter our cognitive lives to the extent that we can so that

¹Frege, Gottlob, *The Thought*; Dummett, Michael, “Truth”.

²Stich, Steven *The Fragmentation of Reason*.

we replace the pursuit of truth with the pursuit of whatever is really valuable for us in the way of belief.

Contrary to Stich I will argue that once we have a clear view of the matter we will find true belief is valuable and that, other things being equal, it is a rational doxastic policy to seek true beliefs. My view is that true beliefs possess both intrinsic and instrumental value but here I will mostly be concerned with the instrumental value of true belief. I will elaborate an argument drawn from some ideas in decision theory which apparently establish both these claims. It turns out that it is possible to measure the utility of various doxastic policies and it also turns out that the utility of this policy is, given certain assumptions, always at least as great as the utility of any *comparable* doxastic policy. I will conclude with a few remarks concerning what features of truth make true belief instrumentally valuable.

It is not immediately obvious how to understand “instrumental value” or how to apply this notion to beliefs. I will say that Q has instrumental value for a person P (at a particular time) if having Q enables P to obtain or increases the probability of P 's obtaining something else she values. On this account the instrumental value of something depends both on the value of what it leads to and the *probability* of its leading to it. A subway token is currently instrumentally valuable (to me) since it can be used to get downtown to a restaurant where I wish to dine. A gamble in which I win a subway token with probability p is also instrumentally valuable. Its value is the expected value of the gamble (the product of the value of the token and p .) Notice that a gamble may have instrumental value even though it actually fails to result in something of value; i.e. you may lose the gamble. An automobile has instrumental value even though owning an it engages one in a gamble in which with a small probability the car crashes. If that unfortunate possibility occurs we would not say that owning the automobile lacked instrumental value.

The probability involved in determining Q 's instrumental value may be either objective or subjective.³ The distinction is important since subjective instrumental value is relevant to evaluating the *rationality* of a policy while objective instrumental value is relevant to evaluating whether or not it is the best possible policy. For example, if X assigns a subjective probability of .6 to a coin's landing heads when the actual probability of its landing heads is .3 then it will be

³Objective probabilities, if there are any are independent of our beliefs. Subjective probabilities are measures of rational degrees of belief.

rational, other things being equal, for X to bet on heads although his best possible policy is to bet on tails.

Beliefs are mental states which provide reasons and are causes of actions. I will assume (with Stich) that a belief is a relation to a representation which possesses truth conditions.⁴ Thus, P believes that snow is white just in case P is appropriately related to a mental representation which has the truth conditions that snow is white. A belief is true iff its truth condition is satisfied.

People don't literally use beliefs to obtain things they value (in the way one uses a subway token) but by acting on a belief one may obtain something of value. So let's say that a belief is instrumentally valuable for P if acting on it is likely to lead to something P values. A property which beliefs may have contributes to the instrumental value of a belief if the belief has its instrumental value (or some of its instrumental value) because it has that property. For example, a round object in my pocket is instrumentally valuable to me it is a subway token not because it is round.

Stich's view is not that true beliefs are *never* instrumentally valuable. It is obvious that sometimes acting on a true belief leads to the satisfaction of a desire. But he does claim that there are plenty of cases in which acting on true beliefs is not instrumentally valuable and he seems to think that even in those cases in which acting on a true belief leads to desirable consequences its doing so is not in virtue of the belief's *being true* but rather in virtue of some other features of the belief. Sometimes he puts his point more cautiously saying that the doxastic policy of seeking true beliefs is not instrumentally *optimal* or even more modestly that it would be difficult to show that seeking true beliefs is generally or even occasionally instrumentally optimal. He concludes that the doxastic policy of seeking true beliefs is not rational or is not or may not be as rational as alternative doxastic policies.

Stich's strategy for attempting to argue for these claims consists of raising three considerations. The first is that sometimes having a true belief can lead to the frustration of a person's desires rather than to their satisfaction. He gives the example of Harry who has a morning flight which he knows departs at either eight or nine. Harry thinks (falsely) that the flight leaves at nine. Actually the flight leaves at eight and is destined to crash. Hearing of the crash after he arrives late to the airport Harry is sure to be thankful that he had the false belief. This example (it is easy to construct others) shows that sometimes having a true belief—the belief that the plane leaves

⁴This account of belief is not essential to the argument.

at eight— can lead to disastrous consequences and that a person may be better off for having a false belief.

Stich observes that to establish the instrumental optimality of truth (and truth seeking doxastic policies) one must show not only that it is instrumentally better than falsity (and false seeking doxastic policies) but better than any of the countlessly many other semantic (and for that matter non-semantic) properties and policies. For example, let “is true*” be a semantic predicate defined in the usual way in terms of reference* and satisfies* where these relations differ from reference and satisfaction in extension (e.g. “Socrates” refers* to Aristotle, “is wise” is satisfied* by athletes etc.). Then the extension of “is true*” may differ from the extension of “is true”. Note that it is only in cases in which “is true*” and “is true” differ in extension that one could have more instrumental value than the other (although even when they have the same extensions it may be that the instrumental value of a belief which has both depends on one property rather than the other).

Stich seems to think that the mere number of such properties makes it implausible that seeking true beliefs is always the optimal policy. Further, it is easy enough to imagine situations in which beliefs possessing one of these properties is more conducive to the satisfaction of a person’s desires than the belief’s being true. For example, there is the semantic property True% which is possessed by the belief that the plane departs at nine (but not by the belief that the plane departs at eight) and possessed by every other belief just in case it is true. Stich points out that Harry would be better off having True% beliefs than True beliefs concerning the time of the fight’s departure. This suggests the following argument. Since there will be some cases in which a person would be better off having a belief which is not true rather than a true belief we can always find a property e.g. True% different from truth which would be better... that is better for the person’s beliefs to have. This seems to show that the policy of seeking true beliefs is not always the optimal doxastic policy.

The third consideration is one that Stich raises in connection with intrinsic value but which he seems to think applies to instrumental value as well. He observes that according to the most plausible account truth is a complex natural property of mental representations whose nature provides little reason to think that its possession by beliefs is of any value. Basically, the account is that e.g. a token of the mental representation “the cat is on the mat” is true iff the relation referred to by the token of “is on” applies to the ordered pair of objects referred to respectively by “the cat” and “the mat”.

The reference relation is characterized in terms of causal relations connecting mental terms with objects and properties. Stich calls this “the causal-functional account of truth and reference”.

No one has succeeded in spelling out the details of the causal-functional account but Stich thinks that if this program proves to be successful the relations will be very complicated ones and, given what we currently know (e.g. about the references of proper names), it is very unlikely that reference and truth will be natural kinds. Further, he points out that the primary criteria for evaluating any proposed account of these relations is whether or not the propositional attitude attributions which an account of reference yields agree with our pre-theoretical “intuitions” concerning these attitudes. That is, an account is tested by seeing whether the beliefs, desires, etc. which it predicts people have coincide with our intuitive verdicts. Stich then points out that our intuitions are malleable and subject to retraining. Presumably he thinks that we can be subject to retraining so that we alter the belief (and other propositional attitude and meaning) attributions we find acceptable. For example, Putnamian retraining may lead us to assigning to someone the belief that H₂O is wet where previously we assigned the belief that H₂O or XYZ is wet. Such retraining would alter the extension of “is true”. Another way of putting the point is that if intuitions changed then the extension of “is true” would be what the extension of “is true*” or “is true**” etc. is. Given all of this he claims that it is implausible that the reference relation and the property truth is very important or that its possession by beliefs is conducive to the satisfaction of the believer’s desires.

Stich draws the following conclusion from these considerations.

... the instrumental value of true belief is far from obvious. It is surely not the case that having true beliefs is always the best doxastic stance in pursuing our goals, and it would be no easy matter to show that believing the truth is generally or (even occasionally!) instrumentally optimal. (124)

When Stich speaks of “instrumental value” it is not clear whether he means subjective or objective instrumental value. We will soon see that the status of his claims depends on this distinction. But whichever he means it is clear that his arguments at most raise suspicions that true beliefs may not possess instrumental value. They fall far short of establishing the claim that “true belief does not have any value...” In fact, the example of Harry does not even establish that acquiring a true belief has less instrumental value than acquiring a true belief about the plane’s departure time. Recall that the

instrumental value of something depends on the *probability* of its leading to something of value. Consider the following two gambles: (G1) if any card other than an ace is drawn (at random) you win \$1,000 and if an ace is drawn you lose \$1,000; (G2) if an ace is drawn you win \$1,000 and if any card other than an ace is drawn you lose \$1,000. Most of us assign probabilities and utilities which make G1 of greater instrumental value than G2. Despite this G2 may actually lead to better consequences than G1 (it does if an ace is drawn). The situation is similar for hapless Harry. His true belief that the plane leaves at eight is like a gamble. With a certain probability it leads to his catching a plane which doesn't crash and with a much smaller probability leads to his catching a plane which does crash. Since the probability of the plane's crashing is tiny this gamble has positive instrumental value *even if the belief leads to his taking a plane which crashes*. So Stich's example neither shows that true belief has no instrumental value or less instrumental value than a false belief.

Although Stich's arguments don't prove that true belief lacks instrumental value or that the doxastic policy of acquiring true beliefs is not instrumentally valuable, the considerations which he raises present a challenge to the philosopher who believes that truth is instrumentally valuable. I will first answer Stich's challenge with respect to subjective instrumental value by showing that a Bayesian decision maker will assign a positive utility to a doxastic policy which seeks her having true beliefs concerning propositions relevant to her decision problems. It turns out that the policy of seeking true belief has a higher subjective utility than any *comparable* doxastic policy.⁵ This will show that, other things being equal, it is rational for a person to pursue the policy of acquiring true beliefs concerning her decision problems. Subsequently I will discuss the prospects for showing that the policy of seeking truth also has objective instrumental value.

A Bayesian decision maker satisfies the following conditions:

- i) Her beliefs possess degrees which satisfy the laws of probability. The degrees of her beliefs are reflected in her willingness to act on them as expressed in principle iv.
- ii) she changes her beliefs in conformity with Bayes' theorem. This is a rule for updating probabilities which says that upon receipt of exactly the information E one's new probability for H should be $P(H/E) = P(H) \times P(E/H)/P(E)$.

⁵This argument was originally sketched in "The Truth Pays" *Synthese* 1981 where I used it to defend William James' assertion that "the truth pays".

- iii) she assigns utilities to consequences (these are degrees of desirability), and
- iv) she selects actions according to Bayes rule which says that one given a choice among actions $a_1 \dots a_n$ one should select the act (or one of the acts) which has the greatest expected utility.⁶

It is implausible that any human being does or can satisfy these conditions all the time. Coherence among ones' probabilities and utilities (which is what satisfaction of i)–iv) is called) is an ideal in the same way that consistency and logical closure among one's beliefs is an ideal (and it entails this latter ideal). These logical norms do not require that we should strive to satisfy them under every circumstance but they do require that if we become aware of failing to satisfy them then rationality requires that we have some excuse (high cost, insufficient time etc.) not to remedy the failure. My argument will be that anyone committed to this ideal is also committed to pursuing true beliefs.

A philosophically interesting, though not widely known, aspect of Bayesian decision theory is its account of experimentation. An experiment relevant to a decision problem is a procedure which yields various outcomes whose probabilities are conditional on the states of nature in the decision problem. It turns out that if the Bayesian conditions listed above are satisfied then experiments can be treated as gambles (whose outcomes are decision problems) and there is a unique extension of an agent's utility ranking satisfying condition (IV) to experiments. This utility assignment specifies the value of the experiment for the decision maker. It can be shown that under certain conditions (described below) any experiment relevant to X's decision problem D has positive utility for X and, other things being equal, the more accurate the experiment the higher is its utility. It is this result which I will use to answer Stich's challenge to demonstrate the value of truth.

First I will briefly explain the Bayesian account of experimentation. Suppose that agent X is faced with the following decision problem:

⁶The above four principles characterizes a standard Bayesian view. There are many variants. On some versions states of belief and desire are represented by sets of probability and utility functions (Levi, *The Enterprise of Knowledge*); on some conditionalization is one among other ways of rationally permissible belief change (van Frassen *Laws and Symmetries*); On some maximizing expected utility is restricted to situations in which probabilities are causally independent of actions. I believe that my argument for the rationality of seeking truth survives in these variants but I leave an investigation of this for the future.

Decision Problem I

	S_1	S_2
d_1	u_{11}	u_{12}
d_2	u_{21}	u_{22}

(The d 's are alternative actions and the S 's are states of nature and the u 's are the utilities which the agent believes will be obtained if she does d when S .)

The value of decision problem I is the expected utility of the act selected by Bayes' rule (i.e. the maximum of $u(d_1) = P(S_1)u_{11} + P(S_2)u_{12}$ and $u(d_2) = P(S_1)u_{21} + P(S_2)u_{22}$). Suppose that X can perform an experiment with outcomes e_1 and e_2 whose probabilities conditional on S_1 and S_2 are $P(e_1/S_1)$ and $P(e_2/S_2)$. The value of acting after experimenting, $u(E)$, is given by

$$u(\text{after } E) = \sum_i P(e_i) \max \left[u \left(\frac{d_1}{e_i} \right), \frac{d_2}{e_i} \right] \tag{U}$$

where $u(d_1/e_i) = \sum P(S_k/e_i)u_{ki}$

$$u(E) = u(\text{after } E) - u(D). \tag{U*}$$

$U(\text{after } E)$ is obtained by treating E as a gamble whose outcomes are the decision problems which would be obtained if the experiment resulted in e_1 or e_2 . These decision problems are identical to I except that the probabilities of S_1 and S_2 are conditional on the outcomes of the experiment. The values of these decision problems are the values of their Bayes' acts. This calculation is illustrated in the examples discussed below.

It is easy to see that $u(E)$ is always greater than 0 and is positive as long as there is a positive probability prior to experimenting that the result of the experiment will alter which act is the Bayes act. Further, and most relevant to our discussion, as $P(e_1/S_1)$ or $P(e_2/S_2)$ or both increase $u(E)$ typically increases and never decreases. That is, the more accurate the experiment the greater is its value. This is strongly suggestive of the claim that acquiring true beliefs relevant to a decision problem prior to acting is valuable to the agent.

How can this result be applied to the question of the value of various doxastic policies? Stich writes as if he thinks that seeking truth is one doxastic policy among many possible ones e.g. seeking true* beliefs and even seeking states of mind which are not semantically evaluable. But can we adopt a doxastic policy other than seeking true beliefs? e.g. by subscribing to *The National Inquirer* instead of *The Times*? There is the following problem. Anyone who has the concept of truth will recognize that if she believes that p she is

committed to believing that p is true and vice versa. So any policy which she recognizes as leading to the belief that p she will also recognize as leading to the belief that p is true. It follows that she will recognize the incoherence of a policy which results in the belief that p when p is not true. So if she knows that a certain policy leads to false beliefs then she will believe its negation. For more complicated doxastic policies (like the policy of seeking true! beliefs discussed below) other suitable adjustments in degrees of belief will be made by a rational person aware of the policy. So it is difficult to take seriously Stich's suggestion that we might opt for doxastic policies which don't aim at true beliefs. Despite this I think it will be instructive to compare the value of doxastic policies. To do this I will suppose that the believer knows what doxastic policy she chooses prior to employing the policy but forgets that she selected this policy after she has acquired beliefs and thinks that she acquired the beliefs through a truth seeking policy. I assume that decision maker X can perform an experiment E and can choose either of the two belief fixation policies concerning what she ends up believing at the conclusion of the experiment:

E_t) if e_i occurs then she believes that e_i occurs.

E_{t^*}) if e_i occurs then she believes that e_j occurs (where $j = i$ and j, i can take values 1, 2).

In other words, E_t fixes a belief concerning the outcome of the experiment iff it is true and E_{t^*} fixes a belief iff it is true*. At the conclusion of the experiment X thinks that she has adopted the first policy in either case.

Given this we can now ask the following question: What are the pre-experiment expected utilities of E_t and E_{t^*} ?

The equations (U) and (U*) supply the answer. It will be useful to go through the calculations in this case for a couple of examples. Consider first the following decision problem:

Decision Problem II

	S	$-S$
d	10	0
d'	0	8
	.6	.4

Suppose that $true^*$ is a property which differs in extension from $true$ and in particular differs with respect to the beliefs that S and $-S$

(whichever of S , $-S$ is true, the other is true^{*}).⁷ The value of this decision problem (i.e. the utility of its Bayes act) = 6. Suppose that before acting X can perform an experiment E_t which will result in her believing whichever of S , $-S$ is true or an experiment E_{t^*} which will result in her believing whichever of S , $-S$ is true^{*}. Since true and true^{*} differ in extension with respect to S and $-S$ and since every statement or its negation (but not both) is true^{*} E_{t^*} results in her believing whichever of S , $-S$ is false.⁸

Using equations (U) and (U^{*}) we can calculate the expected utilities of the two policies:

$$\begin{aligned} u(E_t) &= .6 \times 10 + .4 \times 8 - 6 = 3.2, \\ u(E_{t^*}) &= .6 \times 0 + .4 \times 0 - 6 = -6. \end{aligned}$$

The difference in value of pursuing the truth rather than pursuing the truth^{*} in this case is 9.2. So X has reason to select E_t and E_t is in fact more valuable than E_{t^*} since it is more likely (given

⁷Stich characterizes alternatives to truth in terms of alternative mappings from mental language to the world. For example, there is the mapping I^* which is just like the ordinary one I with the exception that the $I^*(\text{"cat"}) = I(\text{"dog"})$ and $I^*(\text{"dog"}) = I(\text{"cat"})$. Truth^{*} conditions are defined in terms of I^* . So the truth^{*} conditions of "The cat is on the mat" are that the dog is on the mat and a sentence is true^{*} just in case its truth^{*} conditions are satisfied. This excursion through alternative mappings is irrelevant to Stich's argument against the instrumental value true belief. Worse, it can be the source of confusion since one may be misled to interpret Stich as arguing that there is no reason to prefer one's mentaleses sentences having the truth conditions specified by I to the truth conditions specified by I^* . It is not easy to understand exactly what is involved in having preferences concerning the interpretation of our mental sentences. There are three possibilities concerning alternative truth conditions. 1. The truth conditions change but everything else including the causal processes connecting stimuli, intentional states and behaviors remain the same. 2. Truth conditions change and the relations between stimuli and belief or between belief and behavior but not both change in such a way as to maintain the usual rational connections. 3. The truth conditions change and relations between stimuli and beliefs (and other intentional states) also change in such a way as to maintain the usual connections of rationality among them.

⁸Stich repeats a number of times that to show that true belief is instrumentally valuable one needs to show not just that it is more valuable than false belief but that it is more valuable than any alternative true^{*} belief. I am not assuming that true^{*} = false since some true^{*} beliefs may be true. I am assuming that if a belief is true^{*} and not true then it is false. Stich seems to reject this since he seems to hold that we can be in doxastic or doxastic like states which are neither true nor false or even lack truth conditions all together. It is not obvious that there are or can be such states but suppose there are. Then Stich thinks that it may be that by altering ones' cognitive life by adopting a policy which sometimes results in such states one might better achieve the satisfaction one's desires.

X 's probabilities) to lead to the satisfaction of X 's desires in this decision problem. Prior to experimentation X will be certain that E_t will lead to the satisfaction of X 's her desires (as represented in the decision problem) and E_{t^*} will lead to their frustration.

The argument is general. It applies to any decision problem in which the acts differ in expected value and the probabilities of the states of nature are not 1. It can be further generalized to situations in which true and true* do not differ in extension with respect to S but X assigns a positive probability to their differing in extension. Then the value for X of seeking truth* beliefs is the value of seeking truth multiplied by the probability that true and true* agree in extension. Obviously seeking truth will always have greater value than seeking truth*.

So for any property true* if X has a decision problem with states of nature S and $-S$ on which true and true* differ in extension (or X assigns a positive probability to their differing in extension) then X has a reason to seek to form a true belief rather than a true* belief concerning S , $-S$. Of course for decision problems on which true and true* have the same extension it makes no difference whether one pursues the truth or the truth* (i.e. which experiment one selects). But even in these cases it is clear that pursuing the true* is valuable only because it coincides in extension with true. If it had not so coincided then it would have been better (by X 's lights) to pursue the truth. This shows that it is in virtue of the belief's being true that it is valuable to seek it.

It is important to see what the argument so far given does not establish. It does not show that truth in general has any value. It applies only to beliefs concerning the states of nature characterizing one's decision problems and only when seeking the truth will result in learning for certain which state of nature obtains. So far I have not shown that there is reason to value the truth of any other propositions. Also, the argument does not show that one has an *overriding* reason to seek the truth. This of course is to be expected since even if true belief is of instrumental value it may be that seeking the truth about some matter may conflict and be overridden by other values. For example if the cost of E_t was greater than 3.2 it would be better to act without experimenting. If 10 (utility units) was given to X if she conducted E_{t^*} it would be better to aim to select it rather than E_t . But E_{t^*} is more valuable than E_t only when there is some external value or disvalue attached to the experiments. In so far as a choice is being made between seeking truth and seeking true* the former has more value than the latter.

Stich might reply to my argument that while it shows that relative to a given decision problem it is valuable to obtain true information about which state of nature obtains we are usually not in a position to achieve such (even local) omniscience. What if instead of acquiring a belief whose content is that one of the decision problem's states of nature obtains the decision maker is only able to acquire a belief with less informational content but which bears on which state of nature obtains. Let's see how this works for hapless Harry's problem. Suppose that Harry is deciding whether to leave for the airport at 8 or 7 (He needs to leave at 7 to catch the 8 am plane etc.). His decision problem might look like this:

Decision Problem III

	<i>E&C</i>	<i>E&-C</i>	<i>N&C</i>	<i>N&-C</i>
<i>d</i> (7)	0	10	10	10
<i>d</i> (8)	8	8	11	11
	.01	.49	.01	.49

E = Harry's plane leaves at eight, *N* = Harry's plane leaves at nine, and *C* = the eight o'clock plane crashes. Let's suppose that *E&C* is the true state of nature. Stich's point is that in this decision problem believing the falsehood that the plane leaves at 9 has better consequences (for Harry) than does believing the truth that it will leave at 8. If Harry believes the latter he will leave home at 7 and be on the plane which crashes (utility = 0) but if he believes the former he will arrive at the airport too late to board the plane and so will survive (*u* = 8).

The example shows two things. One is that actually possessing a true belief can have disastrous results and the other is that, as a consequence, there are circumstances in which one would actually be better off not seeking to have a true belief about a certain matter (in this case the time of the plane's departure). Indeed, Harry would be better off seeking to have true% beliefs (where the belief that the plane leaves at nine is true% iff it leaves at eight) about the time of departure.

I think that this example is Stich's strongest argument against the instrumental value of truth. But neither of the two points just mentioned show that true belief is not instrumentally valuable or that it is not optimally rational for *X* to seek the truth about the time of the plane's departure.

First, let's compare seeking true belief concerning the time of the plane's departure with seeking true% beliefs.

Let E be an experiment whose whose outcomes are e and n where $P(e/E) = 1$ and $P(n/N) = 1$ and let E_t and $E_{t\%}$ be the following two belief fixation policies:

E_t) if the outcome of the experiment is e then Harry believes Eight and if the outcome is n then Harry believes Nine.

$E_{t\%}$) if the outcome of the experiment is e then Harry believes Nine and if the outcome is n then Harry believes Eight. $E_{t\%}$ fixes true% beliefs concerning the outcome of the experiment.

Using (U) and (U*) we can calculate the utility of these belief fixation polices. E_t is a gamble in which if Harry comes to believe E then he wins a decision problem $D(E)$ just like III except the probabilities of the states of nature are the conditional probabilities on the condition E . These are $P(E\&C/E) = .02$, $P(E\& - C/E) = .98$, $P(N\&C/E) = 0$, $P(N\& - C/E) = 0$. If he learns N he wins a decision problem $D(N)$ just like III except the probabilities are $P(E\&C/N) = 0$, $P(E\& - C/N) = 0$, $P(N\&C/N) = .02$, $P((N\& - C/N) = .98$. The value of $D(E) = 9.8$ and the value of $D(N) = 11$. Finally, the value of E_t is the weighted average of the values of $D(E)$ and $D(N) = P(E) \times 9.8 + P(N) \times 11 = 10.4$. We calculate the value of $E_{t\%}$ by supposing that when $E(N)$ is true Harry comes to believe $N(E)$. If E is true Harry will believe N and so do $d(8)$ whose expected utility he will believe is 11. Since E is true the correct expected utility for $d(8) = 8$ and the correct expected utility of $d(7) = 9.8$. Harry thus suffers a utility loss of 1.8. Similarly if N is true Harry will come to believe E and so do $d(7)$ whose expected utility he will believe is 9.8. The correct expected utilities of $d(7)$ and $d(8)$ are 10 and 11 so Harry suffers a utility loss of 1.2. It follows that the expected utility of $E_{t\%} = 8.9$ and Harry suffers a utility loss of 1.5 if he selects $E_{t\%}$ rather than E_t .

As in problem II the expected utility of the policy of believing the true outcome of the experiment is greater than alternative doxastic policies. This holds even when we know that we are obtaining incomplete (relative to the decision problem) information. Of course, this doesn't mean that when Harry actually obtains the partial information and then acts he will end up with high utility. The example shows this. But, as we have stressed a number of times, the fact that a gamble actually has an unfortunate (for Harry) outcome doesn't show that one ought not to have selected it. Since seeking the truth concerning the experiment's outcome is the policy with the greatest utility it follows that true belief is instrumentally valuable.

However, it is not the case that being true is the instrumentally most valuable property for a belief concerning the outcome of the experiment to possess. Consider the property true! which is possessed by either y or n (statements describing the outcomes of the experiment) just in case acting on that belief results in the best possible consequences (for Harry) given the actual state of nature and possessed by every other statement just in case it is true. If Harry is concerned only about decision problem III then it is certainly of greater utility for him to have true! beliefs than true beliefs. But how can Harry seek true! beliefs? If the only experiment which he has available is E then the best way of his seeking true! beliefs is for him to seek true beliefs. If he seeks a true beliefs concerning the outcome of the experiment then the probability that he will also obtain true! a belief is .99. If Harry can conduct an experiment E' which tests the proposition $E \rightarrow C$ (i.e. has an outcome y iff $E - C$) then the policy of believing the truth about the outcome of this experiment is equal in value to the policy of acquiring a true! belief about the outcome of E . If some deity could somehow bring it about that Harry believes the true! outcome of E then Harry's belief would carry the same information that E' delivers. In view of these considerations it seems correct to say that the policy of seeking true! beliefs concerning E is really not comparable with the policy of seeking truth about E but with the policy of seeking true about E' .

I have established that the policy of seeking true beliefs possesses subjective instrumental value (and as great a value as any comparable doxastic policy). And since subjective instrumental value is relevant to the rationality of an action (at least according to the Bayesian account which we are assuming) we have also shown that, other things being equal, it is rational to seek true beliefs. But what about the *objective* instrumental value of seeking true beliefs? In calculating the objective value of an experiment (and a belief fixation policy) for a decision problem we should employ objective probabilities. Exactly what the objective probabilities of the states of nature in a decision problem are is far from clear so it is difficult to answer the question. But we can say this much. First, if the experiment is one which determines which state of nature obtains (i.e. it yields perfect information) then whatever the prior probabilities of the states of nature the experiment will have a maximum expected utility relative that decision problem and the policy of seeking true belief will have at least as great a utility as any other doxastic policy (neglecting, of course, the cost of the experiment). But the situation with respect to experiments which yield partial information is different. In this case it is possible for the objective utility of an experiment

to be negative. This can easily be seen from the Harry example. If the objective probability that the eight o'clock plane crashes is high and the Bayes act for Harry (calculated using his subjective degrees of belief) is taking the nine o'clock plane) and his ticket is really for the eight o'clock plane then, obviously, finding out the time of the plane's departure will have a negative objective utility. And a doxastic policy which results in fixing true* beliefs concerning the outcome of the experiment may have greater utility than the truth fixing doxastic policy. It may be that Stich's skepticism concerning truth concerns objective utility and that he thinks that it is implausible that the policy seeking true beliefs generally has a high objective utility or, in any case, that it is difficult to show that it has a high objective utility. How can we evaluate these claims?

First, it is certainly logically (and as far as I can see physically) possible that we frequently (or always) are in situations like Harry's in which we would be objectively better off not seeking true beliefs or seeking true* beliefs instead of true beliefs. Is there any reason to think that we are so situated in the world? It may be that Stich thinks that if the causal-referential account of truth is correct then it is plausible that we are often in such situations or that it would be difficult to show that we are not often in such situations. He does argue that truth is not a natural kind and is responsive to our intuitions but does this go any way toward showing that seeking true beliefs has a low objectively instrumental value?

I think not. Money is not a natural kind and whether or not something counts as money is responsive to people's attitudes toward it but I think that no one would say that this shows that money is not valuable. In fact, we have a good deal of reason to believe that seeking truth has objective value. So if Stich really could produce an argument for the claim that if the causal-functional theory is correct then seeking truth would not have high objective value then that would be a good argument against the causal-functional theory (but I don't believe that he has such an argument).

Here is why I think that we have reason to believe that the objective value of seeking true belief is high. First, we know from experience that objectively we have been better off when we have obtained true information that bears on our decision problems than when we act prior to obtaining such information or act after obtaining false information. People who look at *accurate* train schedules prior to leaving for the station are more likely to catch trains and satisfy their desires than those who don't although, of course, every once in a while a train crashes or gets stuck in a tunnel and one may which one had missed the train. And physicians who conduct tests

prior to treatment get better results than those who don't. We have overwhelming inductive evidence that we are objectively better off when we obtain true information relevant to our decision problems.

There are two further reasons to believe that the policy of seeking true beliefs will have a high instrumental value. One is that as long as a person's subjective probabilities are close to the objective probabilities the utility of seeking true beliefs will be high. For example, if Harry's degrees of belief in decision problem III are close to the objective probabilities then there will be little difference between the subjective and objective utilities of his seeking a true belief concerning the plane's departure time. However fallible our doxastic methods might be it does seem that they generally result in degrees of belief which are not far from objective probabilities. And, although Stich correctly argues that natural selection cannot support the claim that our doxastic policies (the ones which have evolved) are optimal it does seem to support the claim that they are pretty good. Needless to say, if humans had typically found themselves in situations like Harry's they wouldn't have survived long enough to have invented airplanes.

Of course the fact that it is generally objectively better to seek true information than not to seek information doesn't mean that there is not some other doxastic policy which would be *objectively* even better than seeking true information. We know that as far as Harry is concerned it is objectively better for him to seek true beliefs concerning the plane's departure time. But is this policy one that Harry has any reason to adopt? Given Harry's degrees of belief it would be irrational for him to adopt this policy. If he were to learn that this was actually the best policy then that information would alter his degrees of belief since this information would be tantamount to his learning that the eight o'clock plane will crash.

I conclude that Stich has not succeeded in showing that seeking truth is not a rational doxastic policy nor that truth is not instrumentally valuable. To the contrary, we have a proof that truth seeking is the optimal rational policy (for a rational being) and that truth has a measureable instrumental value (relative to a particular experiment and decision problem). Finally, we have good reason to believe that truth seeking has objective instrumental value as well.*

*I would like to thank Kati Balog, Mark Deffenbaugh, Jerry Fodor, Gary Gates, Andrew Milne, and Fritz Warfield for comments on previous drafts.