

he requirement resembles that of Annis's theory, then we can expect it to be possible to have knowledge while what is required does not exist. Of course other possible versions of subject contextualism would stand or fall on their own merits.

Perhaps one has to be "in a position to know" that contrary possibilities do not obtain. Depending on how this is understood, either it adds no substantial requirement or it is as mistaken as a requirement of actual knowledge. See note 17 below.

John Hawthorne discusses similar principles about knowledge in Chapter Two of his book, *Knowledge and Lotteries* (Oxford University Press, 2003).

For further discussion of this approach to the justification needed for knowledge, and its impact on external world skepticism, see section IV of Richard Feldman and Earl Conee, "Making Sense of Skepticism" in *Evidentialism: Essays in Epistemology* (Oxford University Press, 2004), pp. 277–306.

Are you "in a position to know" e.g. that it is not the case that you see nothing while two Patagonians perpetrate peccadilloes? Well, when you know according to the current internalist view you have a good enough evidential basis to infer his conclusion, if you have the relevant conceptual and inferential equipment. If that is enough to make you "in a position to know" that such contrary possibilities do not obtain, then you are automatically in such a position when you have the internal justification needed for knowledge. If you are not so equipped, and this keeps you from being "in a position to know" that such contrary possibilities do not obtain, then knowledge does not require you to be in such a position.

A closely related issue is that of whether or not you have to be justified in believing that the contrary possibilities do not obtain, in order to know. There is an extensive helpful discussion of this in Chapter 2 of Peter Klein's *Certainty: A Refutation of Skepticism* (University of Minnesota Press, 1981).

What about the possibility of your being in a Gettier case? Purely internal states never entail that one is not in a Gettier case. But this possibility bears the same relation to knowledge as the skeptical possibilities. If you are in an extraordinary situation where some undefeated consideration in your possession makes for some actual chance that you are in a Gettier case, then you do not know. In the usual circumstance where you possess no such consideration, the mere possibility of a Gettier case does not intuitively conflict with your knowing. For more about this, see section IVB of "Making Sense of Skepticism," op. cit.

TRUTH CONSEQUENTIALISM, WITHHOLDING AND PROPORTIONING BELIEF TO THE EVIDENCE

Michael R. DePaul
University of Notre Dame

I. Introduction

Truth consequentialism is a view regarding epistemic evaluation. The basic idea is that true belief is the fundamental epistemic good, and that cognitive practices are to be evaluated by determining how much of this good we would attain as a consequence of employing those practices.¹ But just as in ethics consequentialism requires a developed account of the good, or moral value, to generate specific evaluations, if truth consequentialism is to produce concrete results, it presumably must do more than simply tell us that true belief is good. It must specify the value of believing a true proposition, and perhaps the values of certain other states as well, e.g., the value of disbelieving a true proposition or the value of believing a false proposition.

Alvin Goldman has recently proposed a theory of epistemic value that has the requisite specificity.² He should be commended for doing so. Many people assert that truth is the epistemic good, but I am not aware of anyone else who has proposed a specific assignment of values. With such an assignment, we can really put truth consequentialism to work—and to the test.³ So I am going to focus on the theory of epistemic value that Goldman has proposed, which he calls veritism, but because his proposal is so straightforward and intuitively plausible, I hope what emerges will be more than a narrow criticism of Goldman's view.

II. Veritistic Value

The first thing to notice is that there is not only an epistemically good thing, truth, that we can attain or fail to attain when we form a belief. When

we form a belief we might also end up with something epistemically bad, falsehood. The pattern is the same as in a typical form of hedonism, which takes pleasure to be the good but recognizes that there is also something bad, pain, that must be taken into account. The overall value of a state of affairs according to such a typical hedonism is not, therefore, equivalent simply to the amount of pleasure it contains. The value of the state of affairs is provided by its utility, which is a kind of sum of the (positive) amount of pleasure it contains and the (negative) amount of pain it contains. Similarly truth consequentialists might be expected to have their eyes not merely on truth, but on an amalgam of the good, truth, and the bad, falsehood. Goldman conforms to this pattern and labels the combined value “veritistic value” or “V-value.” So on Goldman’s view it is actually V-value, not merely truth, that truth consequentialists wish to maximize, and any plausible alternative to Goldman’s value theory would have to define some similar value that somehow aggregates epistemic goodness and badness.

Strictly speaking, it is not truth and falsehood that truth consequentialists take to be epistemically good and bad. It is true *belief* that they take to be the epistemic good and false *belief* that they take to be epistemically bad. There are (at least) two different ways in which we can think about belief. According to the first way, there are three discrete stands one can take with respect to a proposition: one can believe it, disbelieve or reject it, i.e., believe the negation, or withhold belief, i.e., neither believe nor disbelieve the proposition. According to the second way of thinking about belief, beliefs come in a seemingly infinite number of degrees or strengths ranging from absolute certainty to total disbelief. These degrees of belief are sometimes thought of as equivalent to subjective probabilities and represented by points in the interval from zero through one. A normal person would likely believe a simple truth of addition to degree one, and if the person thought about it, he or she would believe the negation of such a truth to degree zero. A person who is very sure that a proposition is true, but not certain, might believe the proposition to something like degree .9. If such a person were to consider the negation of such a proposition, he or she would believe it to a low degree, ideally, to degree .1. In providing his account of veritistic value, Goldman does not choose between these two ways of thinking about beliefs, which he calls the trichotomous scheme and the degree of belief (DB) scheme. He instead describes how to assign veritistic values to the various possible belief states on each scheme.

Goldman employs a question answering model according to which “agent *S*’s belief states (of either the trichotomous or the DB kind) have value or disvalue when they are responses to a question that *interests S* (or, more generally, when other agents are interested in *S*’s knowing the answer).” (Goldman 1999, p. 88) Goldman assigns veritistic value only to beliefs that answer questions that interest the agent, or other relevant parties, because he does not feel that a person’s belief system should be

assigned a low V-value just because the person is ignorant of the answers to many uninteresting questions. Goldman opts for a broad understanding of ‘interesting’ that includes questions about which the agent “has an aroused curiosity or concern” as well as questions that the agent would find interesting “if he/she only thought of them, or considered them” or “if he or she knew certain facts.” (Goldman 1999, p. 95)

Within the trichotomous scheme, Goldman assigns V-values to belief states taken towards a true proposition that answers an interesting question as follows: “If *S* believes the true proposition, the V-value is 1.0. If he *rejects* the true proposition, the V-value is 0. And if he *withholds judgment*, the V-value is .50.” (Goldman 1999, p. 89) On the DB scheme, the V-value of a belief in a true proposition is simply equal to the degree of belief.⁴ Goldman is able to avoid assigning V-values to believing, disbelieving and withholding false propositions because he assumes that when one believes a proposition *P* one disbelieves not *P* and that when one believes *P* to degree *n* one believes not *P* to degree $1-n$.⁵ Hence, when one believes or has a high degree of belief in a false proposition, this will impact the V-value of one’s total cognitive state because one will disbelieve or have a low degree of belief in the true negation of the false proposition.

Goldman is most interested in changes in V-value over time. He usually evaluates practices applied to a range of cases, actual and hypothetical, and takes “the *average* (mean) performance of the practice across those applications as a measure of its V-value.” (Goldman 1999, p. 92) Thus, as he deploys it, truth consequentialism is a kind of epistemic rule consequentialism. Individual acts of belief formation are not evaluated with reference to the amount of veritistic value they produce. Rather, we are to consider how much veritistic value we would secure by following various practices, which we can think of as rules for the formation of beliefs so long as we bear in mind that these may or may not be rules we could self-consciously apply to cases and follow. In some social cases, Goldman is interested in the *aggregate* level of knowledge in a community. The mean V-value for individuals in the community can be used as a measure in such cases, but there are other possibilities. Finally, where appropriate, Goldman allows certain factors not represented in the basic V-value system to influence epistemic evaluation of a practice, for example, the speed of information acquisition or the degree of interest in the questions the practice addresses.

One might initially balk at the values Goldman assigns on the trichotomous scheme. “Why,” one might ask, “should the value of disbelieving a true proposition, or equivalently, believing a false proposition, be zero? This is a bad state to be in, so it would seem more natural to represent its value with a negative number. And why assign withholding a true proposition a positive value of .5? Surely this seems to be a neutral state of affairs, and should accordingly get a value of zero, not a positive value.” There really isn’t any substance to such worries. The numerical values assigned really do

not matter, so long as the relation between the values assigned to the three belief states remains the same. We could use -1 , 0 , $+1$ or 0 , 5 , 10 or -100 , -50 , 0 , or, as Goldman chooses, 0 , $.5$, 1 . Whichever of these systems one uses to evaluate alternative modes of belief formation, one will end up ranking alternative belief forming practices in the same way. The substantive idea being advanced here is that believing a truth is as good, epistemically, as believing a falsehood is bad, epistemically, and withholding is neither good nor bad. Or to put the idea another way, the value of withholding a truth stands midway between the value of believing it and the value of disbelieving it. Goldman captures this idea with the numbers 0 , $.5$, 1 , and this is a perfectly good way of capturing the idea. In addition, I think we would have to grant that the idea these numerical values captures *seems* highly intuitive. Suppose someone were to propose assigning different values, say -1 , 0 and 2 , to disbelieving, withholding and believing a truth respectively. This would be to assert that believing a truth is twice as good as believing a falsehood is bad. Surely if no more were said, such a choice would seem arbitrary. "Why not -1 , 0 , 1.5 or 2.5 or 3 or \dots ?" we would want to ask. We would naturally demand some sort of justification for such an assignment in a way we are not led to demand one of Goldman's. His is the obvious default position, the natural place to begin. If you want to try something different, you are going to need a good reason for doing so.

Nevertheless, as natural and intuitively obvious as Goldman's proposal regarding veritistic value is, I think it runs into problems concerning withholding and apportioning degree of belief to the evidence. The first problem I will describe arises because the assignment of V-value is limited to interesting propositions. Since it does not turn on the most fundamental element of the scheme, it is not, perhaps, terribly significant. I will even be able to suggest what might be an easy fix. But the next two problems, which are related, just might challenge the basic idea behind Goldman's assignment of V-values, the idea I have just argued is so highly natural and intuitively plausible.

III. Uninteresting Propositions

Goldman justifies limiting the assignment of V-values to interesting propositions on the ground that a person should not be penalized for being ignorant of the answers to questions that are of no interest. Let's agree. It does not immediately follow that it should not matter at all, epistemically, what stance a person takes with respect to uninteresting propositions. Yet this is what Goldman suggests when he fails to define V-values for uninteresting propositions.⁶ What follows immediately from the idea that a person should not be penalized for ignorance of the uninteresting is that it does not matter epistemically if a person fails to have *true*

beliefs about uninteresting propositions. Although it is a slightly different matter, perhaps we could also agree that a person should receive no veritistic credit for true beliefs in uninteresting propositions. But what if a person employs cognitive practices that lead him or her to disbelieve many uninteresting true propositions, or equivalently, to form many false beliefs in uninteresting propositions?⁷ Surely such cognitive practices should be counted as epistemically and, one would think, veritistically inferior to practices that produce either true beliefs or no beliefs at all regarding uninteresting propositions. It seems to me that false belief is always epistemically bad, regardless of whether the proposition believed is interesting. Hence, although I can agree, at least for the sake of argument,⁸ that uninteresting propositions should be treated differently from interesting propositions in terms of V-value assignment, I do not think it is correct never to assign a V-value in the case of an uninteresting proposition.⁹

I believe Goldman needs some mechanism for penalizing practices that yield false beliefs in uninteresting propositions. Simply assigning a V-value of 0 to disbelief in uninteresting true propositions might well do the trick.¹⁰ Suppose practices π_W , π_R and π_U perform equally well with respect to the same range of interesting propositions, but π_W produces no beliefs at all with respect to uninteresting propositions, while π_R and π_U frequently produce beliefs regarding uninteresting propositions. The difference between π_R and π_U is that when it comes to uninteresting propositions, π_R produces true beliefs just as reliably as it does with respect to interesting propositions while π_U is terribly unreliable about uninteresting propositions. One way this might happen is if π_R produces beliefs about interesting and uninteresting propositions in the same way, while π_U would have us take wild guesses about uninteresting propositions. If we assign a V-value of 0 to false beliefs in uninteresting propositions, but still define no V-value for true beliefs in or withholdings of uninteresting propositions, the mean V-value for π_U would be driven lower than the mean V-values for π_W and π_R by the additional disbeliefs assigned a V-value of 0 . The total V-values for π_W , π_R and π_U would be equivalent, since the totals were equivalent on Goldman's original assignment, and any additional disbeliefs assigned a value by my proposal would get a value of 0 . But because it is unreliable about uninteresting propositions, π_U would have a higher number of output beliefs with a defined V-value than either π_W or π_R . Hence, the total V-value would be divided by a larger number in the case of π_U , thus yielding the smaller mean.

Notice, in addition, that unless π_R is perfectly reliable with respect to uninteresting propositions, a person employing π_R would be led to disbelieve some uninteresting true propositions. As a result, a person employing π_R would have more beliefs with an assigned V-value than a person employing π_W . Hence the mean V-value for π_R would be lower than the mean V-value for π_W . I think this is a happy result. If we accept the idea that a

person who fails to believe all sorts of uninteresting truths is not epistemically worse off than a person who believes such truths, then there is nothing to be gained by forming such beliefs. But if I am right that it is always epistemically bad to hold a false belief, there is something to be lost by forming beliefs about uninteresting propositions—unless one is infallible when it comes to the uninteresting. Hence, it would seem the best epistemic policy is to withhold uninteresting propositions.

There is more to be said in behalf of withholding uninteresting propositions. Although we often ignore the fact, there is a cost to believing things. At the extreme, there is a cost because we are finite beings and can only have so many beliefs: each belief we form takes up resources, and if we reach the limit of our capacity, every uninteresting belief we hold will prevent us from holding some interesting belief. But one might reasonably think there are significant costs, some of which are properly epistemic, associated with uninteresting beliefs even when we are running far below maximum capacity, as all of us surely are. It is appealing to think of our doxastic systems as analogous, in some ways, to an office filing system or the hard disk on a computer. As I well know, if either is loaded up with useless clutter, all sorts of problems arise. The useless material continually gets in the way, making it difficult, and perhaps impossible, to access and use the important material. The analog for a doxastic system of being unable to locate an important file in a cluttered office would be inability to remember something, which would entail failure to form a belief in an interesting true proposition and the consequent loss of V-value. An analog of the slow processing one experiences with an overly full hard drive might be something like finding it difficult to concentrate on answering an interesting question because one is distracted by recalled trivial facts and incidents.¹¹ And even if such problems could be avoided, it surely is not a good thing to accumulate clutter, be it physical, electronic or mental. If nothing else, any process of information acquisition and management that fosters such accumulation is grossly inelegant. (For goodness sakes, even the IRS lets us destroy records after three years!)

Consider two practices that deal exclusively with uninteresting propositions. The first, π_{GH} , is a good housekeeping practice. If we were to employ it, we would withhold belief in any propositions we consider that are uninteresting now and unlikely ever to become interesting and we would also jettison beliefs once it becomes clear that they will no longer be of any interest. The second practice, π_{PR} , is a pack rat practice. Those who employ it are led to acquire and retain true beliefs in large numbers of uninteresting propositions and to do so in an absolutely reliable way. Because neither of the ways of assigning V-values considered so far assigns values to beliefs in or withholdings of uninteresting true propositions, neither π_{GH} nor π_{PR} will have any direct impact on the total V-values or even the mean V-values of the doxastic systems of those who employ them. Nevertheless, although I

could not say I am certain about this, I do think we would want to say that π_{GH} is superior to π_{PR} . It is not implausible to think that at least part of the superiority is epistemic, and hence that the veritistic evaluation of the practices should capture this difference.

Perhaps we can account for the difference between π_{GH} and π_{PR} by looking at the indirect effects of the practices on the interesting beliefs persons employing them are able to form. As the analogies with cluttered offices and hard drives suggest, a person who does not waste cognitive resources on uninteresting propositions will almost certainly be able to do much better with respect to interesting propositions. Moreover, since π_{PR} involves infallibility regarding uninteresting propositions, it is not a practice any of us could in fact employ. Practices we really could employ, and perhaps some of us actually do employ, that allow for the pack ratish accumulation of uninteresting beliefs will obviously involve the formation and retention of some false beliefs. Hence, usable practices that are similar to π_{PR} will have a direct negative impact on the mean V-values of the beliefs of those who employ them. Perhaps these indirect and direct effects are sufficient to account for the epistemic superiority of good mental housekeeping over the accumulation of mental clutter, and the sense that there is more to be said in favor of the former and against the later arises from a confusion of pragmatic and properly epistemic considerations.

IV. Withholding

I have suggested that it is appropriate to withhold belief in uninteresting propositions. I also believe that withholding is very significant in the case of interesting propositions. It is, in my opinion, very often epistemically best for us to withhold belief. Failure to withhold when the circumstances call for it, as I think they often do, can wreak havoc with a person's system of belief. Just think of gullibility and jumping to conclusions. Both involve failures to withhold. I would not think there is the slightest question about the epistemic evaluation of either practice: both are epistemically bad. Any significant system of epistemic evaluation ought to allow us to criticize such practices.¹² More generally, any such system ought to yield a favorable evaluation of withholding in many situations.

Consider a class *C* of propositions such that each proposition in *C* interests a person *S* and *S* takes the evidence to count equally for and against the proposition's truth.¹³ Let us suppose, for simplicity, that the members of *C* are logically independent. Finally, let's suppose that *S* is an accurate assessor of evidence, where *S* is an accurate assessor of evidence just in case the proportion of true beliefs in the class of propositions *S* takes to be supported by the evidence to degree *n* ($0 \leq n \leq 1$) is equal to *n*.¹⁴ So in this case, since *S* is an accurate assessor of evidence, half the propositions in

C are true and the other half false. What sort of practice would be epistemically best for S to follow in forming propositional attitudes towards the propositions in C? What doxastic practice for situations like this should society teach or otherwise seek to inculcate in individuals? Here are two possibilities: The daring practice, π_D , says, "Let the fact that you take the evidence to be perfectly balanced be damned—go ahead and believe the propositions in C!" The cautious practice, π_C , is to withhold judgment on all the propositions in C. I am very strongly inclined to say that the cautious policy is epistemically better. I expect you agree. Unfortunately, truth consequentialism, at least when Goldman's assignment of V-values is taken to provide the epistemic value theory, cannot endorse π_C in preference to π_D . The reason for this is simple: π_D and π_C have the same mean V-values!

Here is the proof: If S were to employ π_C , S would withhold each proposition in C. S would also withhold the negations of all these propositions.¹⁵ Half of the propositions S would withhold are true, so S will earn a V-value of .50 for each of these propositions. Goldman's assignment does not define a V-value for attitudes taken towards false propositions, so S's withholdings of these propositions would not count one way or the other. Since S would get a V-value of .50 for every proposition that figures in the computation, the mean V-value for employing π_C would be .50. If S were to employ π_D , S would believe each proposition in C and disbelieve the negation of each proposition in C. Half of the propositions in C are true, and S would believe all of these true propositions. Half of the propositions in C are false. The negations of these propositions will be true, and S would disbelieve each of these propositions. S would earn a V-value of 1 for each of the true propositions from C S would believe and a V-value of 0 for each true negation of a proposition from C S would disbelieve. The number of beliefs for which S would earn a V-value of 1 equals the number of disbeliefs for which S would earn a V-value of 0. Again, no V-value is defined for attitudes towards false propositions, so these don't figure in the computation. Hence, the mean V-value for employing π_D would be .50, the same as the mean V-value for π_C .

If we took this sort of V-value assessment seriously, we would have to approve even more daring doxastic practices. On the assumption that S is an accurate assessor of evidence, the *only* time a cautious practice of withholding would not have a mean V-value lower than a daring practice would be when the practice applies to a class of propositions like C, i.e., a class of propositions S takes the evidence to support to degree .5. If S is an accurate assessor of evidence, no matter how little evidence S thinks there is or how nearly balanced S takes it to be between a proposition's being true and it's being false, if S takes the evidence to tilt the scales the tiniest bit one way or the other, the policy of believing or disbelieving in accord with the slightest indication will have a higher mean V-value than a policy of withholding. Of course, the mean V-value of a policy of believing propositions that S takes

to have only very slight support will only be very slightly greater than the mean V-value of withholding such propositions, but it will be greater.

I'm not clever enough mathematically to construct a general argument, but it is simple to convince oneself by working out examples. Here's one. Consider the class C* of propositions that S, who is an accurate assessor of evidence, takes the evidence to support to degree .6. For simplicity, let's suppose that there are 100 propositions in C* and compare practices π_D and π_C , analogous to the practices described above, applied to the propositions in C*. If S were to employ π_C to form beliefs about the proposition in C*, S would withhold all the propositions in C* and all the negations of these propositions. Of the 200 propositions S would withhold, there would be defined V-values only for the 100 that are true (60 from C* and 40 of the negations). These V-values would all be .50, so the mean V-value for π_C would be .50. If S were to follow π_D , S would believe all the propositions in C*. 60 of these beliefs would be true, and S would earn a V-value of 1 for each of these. The other propositions from C* are false, so the beliefs in these propositions would have no defined V-value. S would also disbelieve all the negations of the propositions in C*. 60 of these propositions are false, so the disbeliefs in these would have no defined V-values and would not figure in the computation. But S would also disbelieve 40 true propositions. Each would have a defined V-value of 0, so they do not add to the total V-value secured by employing π_D , but they do figure in the computation of mean V-value, lowering it to .60, which is still greater than the V-value for the cautious policy.

V. Apportioning Degree of Belief

Apportioning degree of belief in a proposition to the degree of evidential support for the proposition can be seen as a kind of analog on the DB scheme of withholding when there is insufficient evidence on the trichotomous scheme. Goldman discusses apportioning degree of belief to the evidence, but does so in the context of criticizing the view that apportioning belief to the evidence is the sole intrinsic epistemic value. (Goldman 2001, pp. 34–37.) He argues that apportioning cannot be the sole value because it is unable to account for the epistemic virtues of evidence gathering. If our evidence for and against some proposition is equally balanced, we could fully attain the value of apportioning belief to the evidence by believing the proposition to degree .5. If we were to go out and acquire very strong evidence in favor of the proposition, e.g., by designing a clever experiment and carefully conducting it, we could again fully attain the value of apportioning belief to the evidence by increasing our degree of belief in the proposition to the requisite degree. The odd thing is that by designing and conducting the experiment and then adjusting our degree of belief accordingly, we would not manage to

improve our epistemic situation at all—if the only epistemic value is apportioning degree of belief to the evidence. Our belief in the proposition would be no better apportioned to the evidence after the experiment than it would have been before. Hence, we are left with no epistemic rationale for gathering evidence. Goldman concludes, correctly in my view, that apportioning degree of belief to the evidence cannot be the fundamental epistemic value.

The question he does not raise, but in my view should, is whether we can provide a rationale for apportioning degree of belief to the evidence on the assumption that veritistic value is the fundamental epistemic value. I'll now try to show that just as we cannot adequately account for the epistemic value of withholding in terms of Goldman's assignment of V-values on the trichotomous scheme, we cannot account for the epistemic value of apportioning degree of belief to the evidence in terms of Goldman's assignment of V-values on the DB scheme.

Let's once again compare a conservative belief forming practice with a daring practice. The conservative practice, π_C , simply has us apportion degree of belief to what we take the evidence to be. The daring policy, π_D , has one give whole hearted assent when one takes the evidence to support the truth of a proposition, even if one takes that support to be very weak. I will argue that assessment of π_C and π_D in terms of V-value favors π_D by working through a couple of examples. I make two simplifying assumptions, as I did above: (i) all the propositions involved are interesting, and (ii) S is an accurate assessor of evidence.

Consider the class of logically independent propositions, C^8 , that S takes the evidence to support to degree .8. To simplify computation, suppose there are 100 propositions in C^8 . Now let's evaluate and compare the performance of π_C and π_D as employed by S with respect to the propositions in C^8 . Since S is an accurate assessor of evidence, 80 of these propositions will be true and 20 false. If S were to employ π_C , S would believe all these propositions to degree .8. S would attain .8 units of V-value for each of the true beliefs, which would produce a total V-value of 64. There is no defined V-value for the beliefs S would form in the 20 false propositions in C^8 . However, on the DB scheme S would believe the negation of each of the propositions in C^8 to degree .2. 20 of these negations would be true, so S's beliefs in them would each have a V-value of .2. This adds an additional V-value of 4. S would also believe each of the 80 false negations of the true propositions in C^8 to degree .2, but once again, these false beliefs would have no defined V-value. So if S were employing π_C , the grand total V-value for S's 100 true beliefs, the beliefs for which a V-value is defined, would be 68. The mean V-value for π_C would be .68.

If S were to follow π_D instead, S would secure a V-value of 1 for each of the 80 true beliefs. But S would secure no additional V-value from beliefs in the true negations of the 20 false beliefs in C^8 , since S would hold each of these beliefs to degree 0. These 20 beliefs would count in the computation of

the mean V-value for π_D , however, since they would have a defined V-value. As in the case of π_C S's beliefs in false propositions do not figure in the computation of the mean V-value because they have no defined V-value. So if S were to employ π_D , S would attain a grand total V-value of 80, and the mean V-value for the practice would be .8. The difference between mean V-values for π_C and π_D in this case comes to .12 in favor of π_D .

Imagine now that S applies π_C and π_D to the class of propositions, C^{51} , that S takes the evidence to support to degree .51. To simplify computation let's again assume there are 100 logically independent propositions in C^{51} . Here the difference in mean V-values for π_C and π_D would be smaller, but it would still support π_D . If S were to follow π_C , S would have 51 true beliefs of degree .51, one for each of the true propositions in C^{51} . The total V-value for these beliefs would be 26.01. S would also believe to degree .49 the negation of each proposition in C^{51} . 49 of these beliefs would be true, yielding a total V-value of 24.01. S's false beliefs again would not figure in the computation. So by employing π_C S would secure a total V-value of 50.02, and the mean V-value for π_C applied to the propositions in C^{51} would be .5002. The mean V-value for π_D would be .51, since S would secure a V-value of 1 for each of the 51 true propositions in C^{51} that S would believe to degree 1 and a V-value of 0 for the 49 true negations of the 49 false propositions in C^{51} that S would believe to degree 0. No V-value is defined for the beliefs S would have in false propositions, 51 to degree 0 and 49 to degree 1, so these beliefs would not figure in the computation. The difference in the mean V-values now would be only .0098, but it still favors π_D .

In my opinion these two cases suffice to establish the pattern. At least if we are accurate assessors of evidence, daring, or perhaps we might better say, reckless belief forming practices that would have us believe propositions to a degree that far outstrips the positive evidential support we have for the propositions turn out to be veritistically superior to the practice of apportioning degree of belief to evidential support. But we almost certainly are not absolutely accurate assessors of evidence. Perhaps apportioning degree of belief to the evidence can be shown to be veritistically superior to reckless practice on the assumption that we are not accurate assessors of evidence.

As it turns out, apportioning is veritistically better than reckless practice if we are so badly mistaken about the evidence that more of the propositions we take the evidence to positively support are false than true. But that is the only situation where apportioning is veritistically better than the reckless strategy. If there are an equal number of true and false propositions in the class of propositions we take the evidence to support to some positive degree, apportioning and the reckless strategy will be veritistically equivalent. But if we are accurate about the evidence in only the minimal sense that there are more true than false propositions in the class of propositions we take the evidence to support to a given positive degree, then apportioning

will be veritistically inferior to the reckless strategy for the propositions in that class. If we think the evidence is strongly indicative of the truth of some set of propositions, and more of them are true than false, but far fewer than our assessment of the evidence would suggest, then the reckless strategy is veritistically superior. And if we take the evidence to be only weakly indicative of the truth of some set of propositions, and many more of the propositions in this set are true than our assessment of the evidence would suggest, the reckless strategy is again veritistically superior. Hence, I do not think a friend of veritistic evaluation will get much traction by challenging the assumption that we are accurate assessors of evidence. Consider just how bizarre the following endorsement of apportioning would be: If you are a very inaccurate assessor of evidence, so much so that when you think the evidence supports some proposition it is in fact more likely to be false than true, then, and only then, it is best, epistemically, for you to apportion your degree of belief to the evidence. But this is the extent of the endorsement veritism can give to apportioning, at least as far as I have been able to determine.

VI. Responses

I can think of two general approaches one might take to try to respond to my criticisms of veritism.¹⁶ (1) One might agree with me that cautious practices such as withholding in the absence of sufficient evidence and proportioning degree of belief to the evidence are epistemically better than doxastic swashbuckling, but deny that I have succeeded in showing reckless doxastic practices are veritistically better than conservative practices. Unless my arguments involve some mistaken presupposition or I have made some sort of stupid arithmetical mistake, one would have to do this by trying to identify veritistic costs and benefits of conservative and daring policies that I have missed. (2) One might accept the veritistic evaluations of cautious and daring policies that I have presented, and deny that there is anything wrong with these evaluations—surprising as it may seem, doxastic swashbuckling is epistemically best. We've made an interesting discovery, not uncovered a ground for objection. In ethics, consequentialists have been far from squeamish about swallowing the unintuitive results of applying their theory, so perhaps we should expect truth consequentialists to have similarly tough palates. Still it would be best if truth consequentialists who opt for this approach could find some way of undermining or explaining away the intuition that withholding and proportioning degree of belief to the evidence are the epistemically best policies in the sorts of cases I have been considering. I can think of two ways in which they might try to do this. (A) They might argue that the intuition involves a different concept of epistemic evaluation from the one truth consequentialism is trying to capture.

(B) They might try to show that the intuition is a manifestation of a certain sort of systematic irrationality to which human beings are subject that has been well documented by cognitive psychologists. I will address these three possible lines of response in the subsections that follow.

1. Hidden Costs and Benefits

To show that withholding can have veritistic benefits that I have missed, one might ask us to consider how we would expect people who have employed the cautious policy to go on with their lives. Persons who employ it will retain a motivation to uncover the truth about many of the propositions they withhold. They will be motivated to continue their inquiries until they find sufficient evidence to decide the truth of the propositions one way or the other. And since such people can be expected to have relatively high standards of evidence, or at least higher than those of the doxastic swashbucklers, when they finally do form beliefs, they are likely to be true. Hence, we can expect that, with sufficient opportunity to continue inquiry and a little luck, such persons will eventually secure additional veritistic benefits in the form of true beliefs. Persons employing the daring policy will miss out on these benefits. Having formed beliefs, they will lack motivation to continue their inquiries regarding the propositions believed. In addition, they will miss the opportunity to correct any mistaken beliefs they might have formed following the daring policy and thereby acquire additional V-value. Hence, while the mean V-value for the daring policy may be greater than or equal to that for the cautious policy *at the point where I stopped my analysis*, we can expect that over time the V-value secured by the cautious policy will slowly improve, while that secured by the daring policy will remain static.

A likely story, but the future certainly need not play out in this way. For example, a person who employed the daring policy would have the time, energy and other resources saved by stopping inquiry to devote to other inquiries which may be expected to yield veritistic benefits. Having formed beliefs, the person will use those beliefs in arguments for other beliefs. The true beliefs thus formed will secure additional veritistic benefits in the form of other true beliefs. Finally, it is not as though the person who follows the daring policy is stuck with her mistakes forever. Proceeding on the basis of beliefs that are false, we would expect that the person would eventually run into some sort of problem, which would provide an occasion for catching the error. And each time she does catch an error, she will secure an additional unit of V-value as compared with the half unit secured by a person who forms a true belief after withholding. Correcting errors has a big veritistic up-side.

We cannot, therefore, count on the cautious policy securing more V-value than a daring policy over the long run, although I suppose it is

possible that it will. Because of the uncertainty, this line of argument cannot provide a very secure defense of truth consequentialism, so let's move on to consider whether it is possible to undermine the intuition that the cautious policies of withholding without evidence and proportioning degree of belief to the evidence are epistemically better than the daring policies. As I've already indicated, I'll consider two ways in which one might try to do this.

2A. Internalist Justification

The first approach questions whether the intuition that the cautious practices are epistemically better than the daring practices employs the right concept of epistemic evaluation.¹⁷ If one thinks in terms of a traditional conception of justification, for example, an internalist, deontological or responsibility conception such as that extensively analyzed by R.M. Chisholm,¹⁸ then it will seem perfectly obvious that in many instances cautious practices such as withholding and apportioning are epistemically best. But it has become increasingly clear that there is more than one concept of epistemic evaluation,¹⁹ and in many cases it is devilishly difficult to keep these concepts straight in one's thinking. Try as one might, it is possible to make intuitive judgments about cases that one takes to be employing one of these concepts when these judgments are in fact colored by a different concept. We cannot just assume that truth consequentialists like Goldman are trying to work out an account of the traditional notion of justification. Indeed, there plenty of indications that they are trying to develop a different, but in their eyes very significant, conception of epistemic evaluation. Thus, they might claim that cautious practices just are not better, in the sense that interests them, than daring practices. The thought that cautious practices are better arises only because one has unwittingly slipped back into thinking in terms of the old-fashioned concept of justification.

It is hard to know how to defend against this kind of charge. My aim has been to remain more or less neutral regarding the exact concept of epistemic evaluation that is in play. I wanted to avoid asking whether Goldman's conception of veritistic value captures some sort of epistemic evaluation that especially interests me, so that I could instead consider whether Goldman's specific version of veritistic evaluation succeeds in producing any reasonable, interesting type of epistemic evaluation. I intentionally set out to avoid the sort of profitless confrontation of intuitions that arises when people are using different concepts of epistemic evaluation without realizing it. And so the point I have been trying to make is something like this: Whatever concept of epistemic evaluation one is using, there is something awfully fishy about that concept if it sanctions reckless doxastic practices of the sort I have considered.

In support of my contention that I have not slipped into using the traditional conception of justification, I might point out that in presenting my examples I never used the term 'justified' or one of its cognates. I did not ask whether beliefs formed as a result of employing one or another practice would count as knowledge, since the answers to that kind of question often turn on whether one thinks the relevant beliefs are justified. It is true that I have characterized the practices I do not think acceptable as daring and reckless. Perhaps I have thereby illicitly drawn upon ideas about epistemic responsibility which are closely associated with the traditional internalist conception of justification. But I do not think so. If withholding is recognized as a distinct propositional attitude and withholding true propositions is assigned a V-value that is not the lowest, presumably withholding is sometimes the best thing to do. I would think that, in essence, my argument regarding the trichotomous scheme turns only on the idea that if withholding is ever the thing to do, one such time is when we accurately take the evidence to count equally for and against a proposition. Similarly, if we are willing to think of belief as coming in degrees rather than as being an all or nothing affair, it seems we should accept that there are some circumstances where it is best to believe to something other than the maximum or minimum degree. My argument regarding the degree of belief scheme is based on the idea that one circumstance where it is obviously best to believe a proposition to some degree greater than .5 but far less than 1 is where we accurately take the evidence to favor the truth of the proposition very weakly. And so I deny that my argument turns on the illicit importation of any notion of epistemic responsibility or traditional conception of justification. But you be the judge.

2B. Prospect Theory

The second strategy for undermining the intuition that a cautious belief forming practice is preferable to a daring practice appeals to prospect theory.²⁰ If we assume that people are rational and that utility accurately captures what they value, then we would predict that when faced with a choice, subjects would choose the option that would maximize their expected utility.²¹ But they do not always make such a choice. Indeed the option subjects take to be the best very often is not the one that maximizes expected utility. It turns out that what seems intuitively best to subjects deviates, and deviates systematically, from what we would calculate to be the most rational choice. Prospect theory is a psychological theory of choice designed by Daniel Kahneman and Amos Tversky (1979) to predict the choices empirical investigation has established that subjects actually make.

There is a way of looking at the intuitive judgments about what is epistemically best on which my arguments are based so that they might be explained—or explained away—in the same way that prospect theory

explains why subjects intuitively choose options that are not rationally best. If my analysis has been correct, we might say that according to Goldman's version of a truth consequentialist framework, when *S* takes the evidence regarding *P* to be equally balanced between *P* and *not P*, the expected V-value, or to be cute, the expected veritility, of believing *P* for *S* is equal to the expected veritility of withholding *P*.²² Hence, on the assumption that *S* is rational and V-value is the only value at issue, one would predict that *S* would be indifferent between believing and withholding. Goldman's epistemological theory accords with this expectation insofar as it says that neither believing nor withholding is preferable in such a case, just as traditional theories of economic rationality would say that, in a similar case where monetary value replaces V-value, *S* would be indifferent between the two options. But my intuitive judgment is that withholding is preferable to believing. And I am pretty sure that I am not alone in thinking that withholding is the preferable choice in such a situation. So we have a case where one choice seems best intuitively, but given certain assumptions, e.g., that Goldman's V-value assignments accurately capture the values at stake, that option can be shown not to be the most rational, since the expected veritility of the two options is equal. I obviously think we should respond by holding fast to the intuitive judgment and reject at least one of the assumptions needed to argue that this judgment is not rational, and I would pick the assumption that Goldman's V-value assignments accurately capture the epistemic values at stake.

Here's how prospect theory might be used to explain away my intuitive preference. Prospect theory is a descriptive theory, not a normative theory. It predicts the choices people actually make; it does not tell us which choice is best in terms of securing what we value. Suppose we are presented a case where the only thing of value at issue is money. Cases of this sort used in psychological research often involve presenting subjects with a choice between various wagers at various odds. In many cases of this type, subjects judge that one choice is best, and this is predicted by prospect theory, but this choice does not maximize expected monetary payoff. In such a case, we might plausibly maintain that the choice that maximizes expected monetary payoff nevertheless is obviously the best or the rational choice, and that the choice subjects actually make is not rational. Thus, one can see prospect theory as charting the way in which actual choices systematically deviate from what is rational. If it could be shown that the intuitive judgment that withholding is better than believing in the case I considered above is exactly the judgment one would predict that a person would make on the basis of prospect theory, one might take this to provide grounds for writing off that intuitive judgment as a manifestation of a systematic, predictable deviation from rationality to which human beings are subject. There would then be no pressure to conclude that Goldman's framework should be altered so that the evaluations it yields capture this intuition.

Prospect theory is complicated. At the risk of oversimplification, I will discuss only two elements, which I think are sufficient for the purposes of this argument. According to prospect theory, the first phase of the choice process involves editing "to organize and reformulate the options so as to simplify subsequent evaluation and choice." (Kahneman and Tversky 1979, p. 274) A crucial step in the editing process involves fixing a neutral reference point with respect to which gains and losses are defined. According to prospect theory, the values that help determine choice are assigned to these gains and losses. The second element of prospect theory we must consider concerns the nature of the function that assigns values to gains and losses. Two features of this function are significant: (i) it is "generally concave for gains and commonly convex for losses" and (ii) it is "steeper for losses than for gains." (Kahneman and Tversky 1979, p. 279) Hence, if we are talking about dollars and cents, because of (i) the value of gaining a certain number of dollars typically will be less than twice the value of gaining half that number of dollars. And because of (ii) the value of a gain of a certain number of dollars typically will not equal the disvalue of the loss of that same number of dollars. The value of gaining that number of dollars will be less than the disvalue of losing that very same number of dollars.

Consider now the "choice" between believing and withholding. There would seem to be two natural ways of setting the reference point: one might set it with reference to believing a falsehood, as is suggested by Goldman's assignment of 0 V-value to this state, or one might set it with reference to withholding. If the reference point is taken to be believing a falsehood, then one might use prospect theory to account for the intuition favoring withholding by pointing to the concavity of the typical value function for gains. Because the function is concave, the value of believing the truth will be less than twice the value of withholding the truth. Hence, one would expect people to choose withholding—and a sure .5 gain in V-value—over a 50/50 chance of a V-value gain of 1. If the neutral reference point is instead set in terms of withholding, then the fact that the value function slopes more steeply for losses than for gains does the work, since it will follow from this that the value of a V-value gain of 1 will be less than the disvalue of a V-value loss of one. Hence, the total value of forming a belief which has a 50/50 chance of being true will be less than .5. Either way the reference point is set, prospect theory predicts that people will tend to favor withholding over believing in the sorts of cases under consideration, just as I have, even though the expected "veritility" of believing is equal to that of withholding. If my intuitive preference for withholding can be adequately explained in this way, it provides no reason to think withholding really is epistemically better than believing. Hence there is no reason for Goldman to alter his scheme for assigning V-values. I suspect a similar argument could be developed with respect to proportioning degree of belief to the evidence.

I am not sure how to respond to this line of argument. I am inclined to dig in my heels on the side of my intuition. But how much sense does that make? I might be doing the same sort of thing as a gambler who continues to think that the slot machine he has been playing has just got to hit soon, even after he has been shown that it works so that that each play is independent. I am not, however, generally incapable of recognizing cases where my intuitions play me false. In a case where bets of relatively small amounts of money are at stake, I am quite confident that the dollar amounts of the various gains and losses correspond to the values that are at stake. Hence when I am shown that the bet that seems best intuitively is not the bet that maximizes expected gain in dollars, I have no problem accepting that my intuition is mistaken. I readily conclude that the best bet is the one that maximizes expected monetary gain. Why then don't I draw the corresponding conclusion just as readily in the epistemic case? Perhaps it is because whether Goldman's V-values accurately capture the epistemic values that are at stake is the very question at issue, while the similar question about monetary values is not open in the case of simple bets. Not having confidence that the numbers correspond to the important values in the epistemic case, I am less inclined to give up intuition and side with the result of a calculation. I just cannot help but think about how strange the evaluations supported by the calculations are. Can we really accept that withholding a proposition is epistemically no better or worse than believing it when the evidence equally supports the proposition and its negation? Can we agree that it is epistemically best to believe a proposition, or believe it to the highest degree, when the evidence favors its being true only to the very slightest degree? The answers to these questions seem very obvious, but they are based on nothing more than intuition. Are such intuitions not of dubious worth in responding to an argument that calls those very intuitions into doubt? Maybe, but maybe not, for in a sense these are not the same intuitions we began with. These are now intuitions that have survived, as strong as ever, in the face of the effort to explain them away. They survive in a way the intuitions about the best bets do not. And so perhaps they should count for something after all.

VII. Conclusion

Let's suppose that my argument has succeeded, that I have shown Goldman's system of V-values does not produce epistemic evaluations we can accept. How important is this result? It might seem that it is a narrow problem with the particular assignment of values Goldman has presented, a problem that can be solved by tinkering a bit with his numerical values. I think not. To avoid the problems I have presented, in effect what one would have to do is assign a value to believing a truth that is less than the

disvalue assigned to disbelieving a truth. But any such choice will seem arbitrary unless some sort of principled rationale can be provided for it, and it is not at all clear what such a rationale might look like. A different assignment can hardly be justified solely on the ground that it avoids the problems I have identified. In addition, it is hard to avoid suspecting that as soon as specific numerical values are chosen it will be easy to show that they give rise to problems of their own. Hence, although it is surely somewhat hasty to make it just yet, it does seem a good bet that truth consequentialism, or at least Goldman's veritism, goes astray simply because it assigns specific values at all. But without a specific assignment of values, it is not clear how truth consequentialism will generate specific evaluations of cognitive practices. And so my hunch is that the problem I have uncovered is not narrow after all.²³

Notes

1. Throughout this paper I use 'cognitive practice', 'belief forming practice' and 'doxastic practice' interchangeably. I sometimes use the term 'policy' rather than 'practice', but this variation is stylistic as well; I make no distinction between practices and policies. My use of either 'practice' or 'policy' might suggest to some readers methods of forming beliefs that we choose to employ and self-consciously follow. This is not what I have in mind. For one thing, I do not want to commit myself to our having voluntary control over all or even some of our beliefs. I do not mean to commit myself to thinking that we have voluntary control, at least of any very direct kind, over the ways in which we go about forming beliefs. I definitely want to count as practices or policies ways of forming beliefs that we consciously monitor, but I also want to include ways of forming beliefs that we normally are not conscious of but can become conscious of, and even ways of forming beliefs that we cannot bring to consciousness. I also want to count as practices or policies ways of forming beliefs that we have in some sense chosen and trained ourselves to employ, but also ways that others have inculcated in us, ways that we have just picked up in the courses of our lives, and ways that we were born with or automatically acquired in the normal course of human development.
2. Goldman 1999 develops the proposal at length in Chapter 3 and then applies it to the evaluation of various social belief forming practices in the subsequent chapters. Goldman 2001 does not provide as detailed a presentation of the account of epistemic value, but the basic view figures importantly in the argument for the thematic unity of the epistemic virtues.
3. Notice that one can formulate a version of truth consequentialism that does not seem to require specific assignments of values. One could, e.g., claim that the best cognitive practice is the one that maximizes the number of true beliefs. However, it seems to me that such proposals often involve implicit views regarding the values of various cognitive states. For example, one obvious thing that seems to be presupposed by a truth consequentialism that looks

- only to maximize the number of true belief is that false beliefs have no disvalue. If this were not so, the practice that maximizes true belief might not be as good as some other practice if the maximizing practice also produces many false beliefs while the other practice produces nearly as many true beliefs but no false beliefs.
4. Goldman cautions, however, that there is no easy way of translating between the trichotomous and DB schemes. Thus, e.g., one should not be misled into equating withholding a truth, which gets a V-value of 0.5, and believing a truth to degree .5, which also gets a V-value of 0.5.
 5. Another simplifying assumption Goldman makes is bivalence.
 6. I may be oversimplifying Goldman's view here. He explicitly allows that truth is not the only valuable thing, and even that truth is not the only value that is relevant to the assessment of cognitive practices. But he does claim that veritistic evaluation is the special job of epistemology. It is not, therefore, entirely clear whether Goldman holds that veritistic evaluation, and specifically evaluation with respect to V-values, can capture all epistemic evaluation. It is clear, however, that Goldman holds veritistic evaluation to be the most central and fundamental sort of evaluation for epistemology. And I think this is enough for the argument I am about to offer, since the difference between belief forming practices I will point out is the sort that should underwrite a difference in the most central kind of epistemic evaluation.
 7. For stylistic reasons, I will often formulate points in terms of 'false belief.' This should usually be interpreted to mean "disbelief of a true proposition," a state to which Goldman assigns a V-value, rather than "belief of a false proposition," a state to which Goldman assigns no V-value.
 8. I can only agree for the sake of the argument because I do not have a settled view regarding the treatment of uninteresting propositions. I find Goldman's idea that a person's system of belief is not worse for failing to contain all sorts of uninteresting truths plausible. But on the other hand, I am also drawn to the thought that any well formed true belief or item of knowledge is epistemically good, in some cases as good as possible, regardless of whether the proposition believed is interesting.
 9. On the assumption that any proposition a person believes or disbelieves interests that person to some, perhaps very small, degree, there is no real problem here. It is not as though this assumption has no initial plausibility. After all, if the proposition did not interest the person at all, why would he or she go to the trouble of believing or disbelieving it? I know of nothing, however, at least in Goldman 1999, that indicates he makes this assumption. And I do not think the assumption stands up to reflection. In the course of a normal day, most of us take note of and cannot help but form beliefs about many completely trivial, uninteresting things. (Thankfully, most of these beliefs are very short lived.) For example, I might be looking for someone's number in the phone book. In the course of doing so, I notice all sorts of names. I see that there is a George Malone and a Harry Malone in the Dowagiac phone book as I run my finger down to the number I want, for Sean Malone. I see no reason to deny that I would believe that there is a George Malone in the phone book, at least for a minute or two after I read the name. But surely the proposition does not thereby

- become interesting to me. And similarly for all the other trivial bits of information we pick up and just as quickly discard as we go through the day.
10. In the argument I am about to offer, and the other arguments I offer in this paper, I work in terms of the mean V-values of belief forming practices. I simply consider the number of true and false beliefs produced by the practices. Goldman 1999 prefers analysis in terms of change in V-value. I do not believe that working things out in terms of V-value change would alter the conclusions I reach; it would simply complicate the arguments. I would add to the specifications of the practices I consider that one begin withholding the propositions in question, and then proceed as the practice directs.
 11. Goldman allows that speed of processing can be taken into account in the veritistic analysis of cognitive practices, so even if such distractions do not prevent one from forming true beliefs about what interests one in the end, the delay is a properly epistemic cost.
 12. I am not here implying that gullibility cannot be criticized on veritistic grounds. It can, at least in certain circumstances. As will become clear, I will be trying to show that a certain sort of jumping to conclusions ends up getting a positive veritistic evaluation.
 13. I do not here mean to be picky about what counts as evidence, e.g., by limiting it to other beliefs. Hence, such things as ordinary visual beliefs which we form quite automatically in response to our experience, and do not form on the basis of propositional evidence or reasons, could be members of C, since for my purposes here, I would simply count the visual experience as evidence. I would say the same thing about various memory beliefs, where I would count the experience of remembering as evidence. And similarly for introspective beliefs, intuitive beliefs, and beliefs about the positions of our own bodies.
 14. I want it to be perfectly clear that I am introducing the concept of accurate assessment of evidence and defining it in a purely stipulative way solely for the purposes of this argument. I am not out to capture any intuitive, commonsense notion of accurate evidence or accurate assessment of evidence, nor do I care whether the concept I here define corresponds or fails to correspond with the technical concepts epistemologists, philosophers of science, cognitive scientists, or any other community of specialists might have of accurate evidence or the accurate assessment of evidence. Perhaps it would have been better to avoid any possible prior associations and consequent misunderstandings by designating the property I here define in some arbitrary way, e.g., as property Q, but I decided against that because the presentation seemed to read more smoothly as it is.
 15. Goldman explicitly states he assumes that persons who believe a proposition *P* automatically disbelieve the negation of that proposition, *not P*. (Goldman 1999, p. 88) I assume he makes the corresponding assumption about withholding, i.e., that persons who withhold *P* also withhold *not P*.
 16. I'll ignore a third strategy: Attempting to identify some third kind of doxastic practice that seems to be epistemically superior to both the conservative and daring practices I have considered and that can be shown to be veritistically superior to both practices. While I admit I would be very interested to hear of such a third practice, I will leave it to the friends of veriticism to try to discover one. My excuse for this, aside from sheer laziness and the lack of any good ideas

about where to look for such a practice, is that the discovery of such a practice would not get veritism completely off the hook, since veritism could still be charged with producing the wrong ranking of the practices. Of course, I would have to admit that if veritism succeeded in identifying the best practice, many would find mistakes in rankings further down the list easy to overlook.

17. Goldman said some things in correspondence responding to an initial formulation of my concerns about withholding that suggested this line of argument to me.
18. See Chisholm 1989 for one of many places where he works with this concept.
19. See Alston 1993 and DePaul 2000.
20. When I suggested this approach to Goldman in correspondence, he was not inclined to take it. But I think it represents an interesting line of argument that might have broad implications for epistemology, so I will present it.
21. I am obviously assuming that a certain instrumentalist account of practical rationality is correct for the sake of this argument.
22. At the risk of belaboring the obvious: S takes the probability of P being true to be .5. So from S's perspective, if S withholds P , S has a .5 probability of a V-value payoff of .5 and a .5 probability of a V-value payoff of .5 for an expected veritility of .5. If S believes P , S has a .5 chance of a payoff of 1 and a .5 chance of a payoff of 0 for an expected veritility of .5.
23. I have had the benefit of discussing an earlier version of this paper with the Philosophy Department at Michigan State University. I have also profited from the opportunity to consider the excellent comments made on an earlier draft by Stephen Grimm and Christian Miller.

References

- Alston, W. (1993). "Epistemic Desiderata," *Philosophy and Phenomenological Research*.
- Chisholm, R. M. (1989). *Theory of Knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- DePaul, M. (2000). "Value Monism in Epistemology," Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press.
- Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- . (2001). "The Unity of the Epistemic Virtues," A. Fairweather and L. Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Oxford: Oxford University Press.
- Kahneman, D., and A. Tversky. (1979). "Prospect theory: an analysis of decision under risk," *Econometrica*.

TRUE ENOUGH*

Catherine Z. Elgin
Harvard University

Epistemology valorizes truth. Sometimes practical, or prudential, or political reasons convince us to accept a known falsehood, but most epistemologists deny that we can have cognitively good reasons to do so. Our overriding cognitive objective, they maintain, is the truth, preferably the whole truth and definitely nothing but the truth (Goldman 1999, p. 5; Lehrer 1986, p. Bonjour 1985, p. 9). If they are right, then at least insofar as our ends are cognitive, we should accept only what we consider true, take pains to insure that the claims we accept are in fact true, and promptly repudiate any previously accepted claims upon learning that they are false. I suggest, however, that the relation between truth and epistemic acceptability is both more tenuous and more circuitous than is standardly supposed. Sometimes I contend, it is epistemically responsible to prescind from truth to achieve more global cognitive ends.

At first blush, this looks mad. To retain a falsehood merely because it has epistemologically attractive features seems the height of cognitive irresponsibility. Allegations of intellectual dishonesty, wishful thinking, false consciousness, or worse immediately leap to mind. But science routinely transgresses the boundary between truth and falsehood. It smoothes curves and ignores outliers. It develops and deploys simplified models that diverge sometimes considerably, from the phenomena they purport to represent. Even the best scientific theories are not true. Not only are they plagued with anomalies and outstanding problems, but where they are successful, they rely on laws, models, idealizations and approximations that diverge from the truth. Truth-centered epistemology, or *veritism*, as Alvin Goldman calls it, easily accommodates anomalies and outstanding problems, since they are readily construed as defects. The problem comes with the laws, models, idealizations, and approximations which are acknowledged not to be true but which are nonetheless critical to, indeed constitutive of, the understand-