

# Utilitarian Epistemology

Stephen Petersen

Niagara University

steve@stevepetersen.net

August 22, 2008\*

**Author:** Stephen Petersen

**Institution:** Niagara University

**Email:** steve@stevepetersen.net

**Address:**

Department of Philosophy

PO Box 2043

Niagara University NY 14109

**Paper word count:** 5,531

---

\*This is a *draft* (version: *utilitarian-epistemology.tex,v1.13*), and should not be circulated or cited without permission.

## UTILITARIAN EPISTEMOLOGY

Standard epistemology has it that there is a particularly epistemic type of value. A belief might be disastrous in many other ways; it might bring about great misery, and be horribly ugly to contemplate, and even cause you to forget many other important things. If that belief is an instance of knowledge, though, then it is valuable in at least one important way—epistemically.<sup>1</sup>

That such special epistemic value exists is a normative claim, and it had better be consonant with one's overall theory of value. If one is a Kantian about value, it is easy to make room for an epistemic dimension; epistemic value is (roughly) from forming correct intentions with respect to believing. It is also easy to accommodate epistemic value if one is an Aristotelian; epistemic value inheres in stably virtuous epistemic habits. Both such pictures of epistemic value have dedicated proponents.<sup>2</sup> According to what we might call “classical” act-utilitarianism, though, the *only* thing of value is welfare, and individual acts are only good instrumentally, insofar as they contribute to it. Thus it seems impossible for such a utilitarian to account for particularly epistemic value. If some belief brings about more misery than happiness, it is of disvalue, however true, justified, *etc.* it may be.

Nonetheless, I think the utilitarian can do epistemology. In fact, I think a utilitarian account of epistemic value provides a powerful approach to puzzles that arise on the other two pictures. I sketch such an account and its potential implications here.

### 1 Utilitarian epistemology

In a certain strict sense, the classical utilitarian must say that there is no epistemic value; only states of welfare are valuable or not. Knowledge is not good in-itself, and neither is understanding,

---

<sup>1</sup>Standard epistemology holds that knowledge is what's of typically epistemic value; Kvanvig (2003) and others have argued otherwise, but they still agree that there is particularly epistemic value in something, such as understanding.

<sup>2</sup>See *e.g.* Chisholm (1977) or Ginet (1975) for the former, and Goldman (2000) or Sosa (1997) for the latter.

<i>epistemology</i>	<i>finance</i>
true belief	profit
knowledge	earnings
mere true belief	windfall

Table 1: The financial analogy

nor epistemic justification, nor even true belief. All are good only insofar as they *tend* to enhance welfare. On this picture, then, the normal objects of epistemic evaluation are good only in the same way that money is good. Money, all agree, has neither intrinsic nor final value. Still, financiers study it and its means of acquisition a great deal, because it is typically of such instrumental value. Epistemology, on the utilitarian picture, is just like the study of finance.

Let us develop this analogy, since I think it runs fairly deep. Think of true beliefs as analogous to money, so the gaining of a true belief is analogous to profit. Building on the common idea that knowledge is true belief that was not “lucky”,<sup>3</sup> we can say knowledge is like *earnings*—that is, money gained through sound investments. Mere true belief, on the other hand, is like a financial *windfall*—money gained by luck.

Now, think again of the piece of knowledge that brings about great net misery. The utilitarian is committed to saying that this belief is of disvalue, and thus it seems that the utilitarian cannot account for epistemic evaluation. But consider the analogous circumstance: someone’s financial earnings bring about great net misery. This possibility does not show that a utilitarian cannot study finance. We might even say those earnings were nonetheless of “financial value”, but this does not contradict anything in the utilitarian dogma, for it is clear to us in this case that to be of “financial value” is a purely *descriptive* claim; it means something close to “makes money”. The utilitarian epistemologist hears “epistemic value” similarly; we can say that knowledge (and more generally true belief) is of epistemic value without thereby committing to the existence of a value independent of welfare.

---

<sup>3</sup>Duncan Pritchard takes this to be a consensus and a “platitude” (Pritchard, 2005, p. 1), citing cases like that in Steup (2006); at any rate, whether platitudinous or not, it is at least a good start.

Why not say “of financial value” means *exactly* “makes money”—and by analogy that “of epistemic value” means exactly “gets true beliefs”? Well, on the financial side, consider an instrument where you can risk all your assets for a 1% chance at a 1,000-fold return. Looking strictly at making money this is a good bet, but I think the right *financial* advice is still to abstain, and I think there are analogous cases on the epistemic side. Perhaps the best characterization of “[financial | epistemic] value” is “[money | belief]-related kind of thing that *tends* to get utility”.

Note that this does not make the account rule-utilitarian, since these are still descriptive claims; the right thing to do is still the utility-maximizing thing, even if it does not get [financial | epistemic] value.

Note also that this talk of “tendencies” is essential to doing utilitarian normative ethics generally. When we look at a *type* of thing, rather than a very specific thing in a very specific circumstance, we cannot talk about its actual utility; presumably anything with causal powers could in principle lead to low welfare or high. To do normative ethics requires making some generalizations, though, and so utilitarians resort to averages and expected utility. Thus when a classical act utilitarian says a charitable type of act is good and a murderous type of act bad, they mean that charitable acts have a high *expected* utility (absent further information) and murderous acts have a low one. The point is, for the utilitarian to generalize about the value of anything (other than welfare) requires a certain degree of uncertainty about particulars.

One special case of such utility-based generalizations is the financier, who can say that buying shares in a diversified mutual fund is a good kind of financial act, and playing high-stakes roulette at the local casino is a bad type of financial act—even though a) the mutual fund might tank and the roulette might pay off, *and* even though b) money from a profitable mutual fund might finance Halliburton’s dastardly schemes, while money lost at the casino might ultimately save starving children. According to my proposal, epistemology is another special case of such utility-based generalizations. The utilitarian epistemologist can say that deduction is a good type of cognitive act, and wishful thinking a bad one—even though a) deduction might get a false belief and wishful

thinking a true one, *and* even though b) a deduced true belief might bring misery and a wishfully thought false belief might alleviate it.

So you might agree that in this sense, at least, the utilitarian can rescue a notion of epistemic value. You might be left wondering, though, about the point of such a rescue. Shouldn't the epistemologist simply study epistemic value, and leave questions about the nature of that value to the ethicists? Predictably, my answer here is “no”, for many central epistemic questions are inextricably bound with ethical ones. That is because, at its core, epistemology is the study of *good thinking*.<sup>4</sup> I'll demonstrate the connections by outlining the implications of utilitarian value for just two current issues in epistemology: the “value problem” for knowledge, and the relation between epistemic justification and epistemic responsibility.

## 2 The value of knowledge

One obvious place of intersection between value theory and epistemology is in the comparatively recent resurrection of a problem from the *Meno*: namely, that of explaining why (or whether!) knowledge is more valuable than mere true belief. After all, as Plato pointed out, both knowledge about the location of Larissa and mere true belief about its location will get you to Larissa.

### 2.1 Current answers

Here is a very brief and somewhat tendentious overview of the current state of play with respect to this “value problem”. It is indebted in many places to the nice exposition of Pritchard (2007b).

Standard process reliabilists have trouble answering Plato's question, since they claim the difference between knowledge and mere true belief is in whether the belief was produced by a reliable

---

<sup>4</sup>Perhaps you protest that epistemology is really the study of knowledge. Yes, that has etymology on its side; but *if* knowledge is the one defining characteristic of good thinking, then my characterization subsumes it. If not though—if knowledge is only part of good thinking, or none at all—my characterization is ecumenical enough for these possibilities too. Surely disparagers of knowledge such as Williams, Stich, and Kvanvig are epistemologists?

process or not. Roughly speaking this means that knowledge and mere true belief differ only in extrinsic properties, so—assuming that the final value of a state supervenes on intrinsic properties only—there can be no value difference between the two. Linda Zagzebski famously compares the situation to two good cups of espresso, one produced by a reliable machine and the other produced by an inconsistent machine that happened to work well this time.<sup>5</sup> The fact that one came from a reliable device does not give any more value to the espresso, and reliabilists seem committed to an analogous position on knowledge. Put this way the puzzle is sometimes called the “swamping problem”, since it seems the value of the product (good espresso / true belief) “swamps” the value of the producer (machine that normally makes good espresso / process that normally makes true beliefs); when you have the former, the latter is of no additional help.

The reliabilist might insist that the knowledge state consists in part of the production process. Thus if the process is valuable (the way a reliable espresso machine is itself valuable), the knowledge state inherits that value by having the process as a part. It also gets value from its true belief part, and so the total value is greater than the value of a mere true belief. First of all, though, this wide construal of the knowledge state is tricky to spell out in detail. Second of all, it still runs up against what Duncan Pritchard calls the “secondary value problem”: that of why knowledge is more valuable than any of its proper parts. An easy way to see this is to consider Gettiered beliefs—they would still have the value of the generally reliable process, and of the (accidentally) true belief, but intuitively they are not just as valuable as knowledge.<sup>6</sup>

So reliabilism seems to be in some trouble here. It’s not noted as often, though, that deontological internalists, too, have their own work to do explaining the value of knowledge. They can explain easily enough how value attaches to *good reasoning*, via dutiful epistemic decisions. But if we see this reasoning as a process distinct from the resulting knowledge, we can still wonder why the result has value in virtue of having so arisen. And even if we have a story here about how

---

<sup>5</sup>Zagzebski (2003).

<sup>6</sup>The obligatory citation: Gettier (1963).

the reasoning process is an intrinsic part of the knowledge state, we are again left with the “secondary value problem”. The deontologist can say that (intrinsic) value attaches to good reasoning, and (instrumental) value attaches to true belief, but you could have both of these without having knowledge, so the value of knowledge over Gettiered true belief remains a mystery.

The virtue epistemologist’s “credit” solution has perhaps made the most progress toward an answer to the value problem.<sup>7</sup> Their first key move: to answer Gettier problems, it is not enough to have the cognitive virtue and the resultant true belief; to be knowledge requires that one have the true belief *because* of the cognitive virtue. Their second key move: extrinsic value does not mean merely instrumental value, and so final value does not supervene on intrinsic properties. That is, though knowledge may derive its value from its relationship to a cognitive virtue, it may still be of final value rather than instrumental. The most plausible analogy here is to that of an Olympic gold medal—whether it was earned at the Olympics or bought on eBay makes a difference to the final value of that medal.<sup>8</sup> Similarly, the virtue epistemologist claims, knowledge is true belief resulting from exercising a cognitive virtue, and so represents a cognitive achievement on the part of its owner. Knowledge, but not mere true belief, is of credit to the believer, and thus gains extra value.

This account too has serious problems, though. First, as many have pointed out, it is not clear that knowledge always represents an achievement on the part of the believer. We can gain genuine knowledge through perception or testimony, it seems, yet it’s hard to claim that such unreflective belief mechanisms are of credit to their owner. Second, it seems to me the standard examples of relational-but-final value are not clearly analogous with knowledge. The most common example is that of a dress once worn by Princess Diana—it gains final value through its relation to its former owner. Both in this case and the medal case, though, the value at hand is symbolic, and nostalgic. The thing is of value because it stands for and calls to mind something else of value. It would

---

<sup>7</sup>Greco (2003); Riggs (2002); Sosa (2003); Zagzebski (2003), *et al.*

<sup>8</sup>Riggs (2002).

be strange if the value of knowledge were similarly symbolic, I think—as though I valued my knowledge because I could reminisce upon it fondly, and bask in the memory of my noble efforts in its acquisition.

## 2.2 Utilitarian epistemology and the value of knowledge

How, then, does the imagined utilitarian epistemologist approach this problem? If the value problem is understood to be the question of why knowledge has more *final* value than mere true belief, the utilitarian’s answer is immediate, but revisionary: knowledge does not have more final value than mere true belief, because both have exactly zero. To the classical utilitarian under consideration only states of welfare are of value; all other value is instrumental.<sup>9</sup>

We might still ask, of course, why knowledge has more *instrumental* value than mere true belief, and this would be closer to the utilitarian spirit. There is also a tradition behind this understanding of the question; in fact it seems to be how Plato heard it, since he answers his question with the assertion “knowledge exceeds true belief by its binding ties.”<sup>10</sup> Knowledge, Plato seems to suggest, is true belief that has been tied down so that it can’t easily escape. It is thus more valuable to us in the same way a valuable statue in an alarmed display case is more valuable than that same statue left unguarded.

This answer is not satisfactory, though; as Timothy Williamson aptly put it, “surely [Plato] recognized that mere true beliefs can be held with dogmatic confidence, and knowledge lost through forgetting.”<sup>11</sup> That is, Plato’s answer does not explain why knowledge is more valuable than a lucky true belief held dogmatically, since both are apt to stick around. It also fails to explain why knowledge we are about to forget is more valuable than mere true belief we are about to forget, given that both are about to disappear.

Williamson amends this answer to one Platonic in spirit; in summary, Williamson says knowl-

---

<sup>9</sup>There are complications here about whether knowledge can help *constitute* welfare; I here assume it cannot.

<sup>10</sup>The Greek here for “by its binding ties” is ‘δεσμεῖ’ (Plato, BCE, 98a).

<sup>11</sup>Williamson (2000) p. 78.

edge is more (instrumentally!) valuable because the justification that comes with knowledge makes it more probable that the knower will keep the true belief when new doubts arise. For example, suppose the mere true belief is a Gettiered one based on false premises. Knowledge is more valuable than that, on this story, because the Gettiered belief is more vulnerable to having its justification undermined. John Kvanvig argues that Williamson's answer is also unsatisfactory, however, for in many ways knowledge is significantly more fragile than the corresponding mere true belief. You merely need a defeater to lose the former—perhaps just thinking of skeptical scenarios will do!—while something must change your mind to lose the latter. *Perhaps* the world is such that you're on balance more likely to keep knowledge than mere true belief, but this would be at best a contingent thesis, and so could hardly capture the intuition that it is in the nature of knowledge to be more valuable than mere true belief.<sup>12</sup>

I think Kvanvig's point here generalizes; any account that explains how knowledge is better than mere true belief by adverting to instrumental value is going to be hostage to contingencies. Any piece of knowledge could, in the right circumstances, be far less valuable than the corresponding mere true belief, and even generalizations about how knowledge is *usually* more valuable will depend on the local causal tendencies. It looks like instrumental approaches to explaining knowledge's value will at best result in empirical hypotheses. Such answers do not seem, on reflection, to be in the spirit of the question.

Prospects for explaining knowledge's value are even worse when we consider the utilitarian analogy developed earlier. To ask "why is knowledge of more instrumental value than mere true belief?" is to the utilitarian like asking "why are earned profits of more instrumental value than monetary windfalls?" The answer to the financial version of this question is clearly that the earnings are *not* more valuable. By analogy, then, neither is knowledge. The epistemic utilitarian embraces this conclusion and denies the intuition that knowledge is better than mere true belief, even on the instrumental version of the value question. The reason, illustrated by the analogy,

---

<sup>12</sup>Kvanvig (2003) chapter 1.

is fairly simple: like anything but welfare, epistemic states are at best of instrumental value, and *generalizations about instrumental value only make sense under uncertainty*. Generally charity is more valuable than murder, but to the utilitarian (and to the utilitarian alone) it is not sensible to ask “why is a charitable act more valuable than a murder that results in the same amount of utility?” To assume there is an answer here begs the question against the utilitarian. The same goes, one step removed, for the question “why are earnings more valuable than windfalls?” Investments with high expected value are in an important sense more valuable than those with poor expected value, but this question builds in the assumption that both result in the same instrumental value (given of course that all else is equal). Finally, the same goes for knowledge and lucky-but-true belief; in the description of the case, both have gotten the relevant epistemic (instrumental) good. To stipulate that, despite the odds, luck-sensitive belief formation nonetheless resulted in a true belief is just like stipulating that the murder under consideration ended up net benefitting people, or that the stupid casino bet ended up paying off. In summary, to return to still another analogy: the epistemic utilitarian has justification for saying the situation posited *is* like Zagzebski’s two cups of espresso. Absent information about how they turned out, we should take the espresso from the reliable machine, of course. But if we learn both cups ended up equally tasty, then we should be indifferent.

Unlike in the espresso case, though, we do have the stubborn intuition that knowledge is more valuable than mere true belief, and the utilitarian needs to accommodate this intuition in some way. Such a task is familiar ground for the utilitarian, though; she already has to explain away intuitions about the theoretical possibility of justified torture and such. So she can simply borrow the typical response strategy already available for such cases. It runs roughly like this: because we are finite creatures who cannot be arbitrarily sensitive to the infinite vagaries of each situation, we rely on heuristics to guide us—rules of thumb about what generally is of use.<sup>13</sup> These become ingrained in us so deeply that we are tempted by deontological intuitions to the effect that an act can be

---

<sup>13</sup>See Harman (1977) p. 155.

wrong even if it makes the world better. Consider this offer, for example: you can choose between two lives of exactly equal welfare, but in only one of those will you also be filthy rich. We might feel an intuitive pull to pick the rich life, even though it has been stipulated that they are worth the same, at least as far as enjoyment of one's own life goes. Similarly, perhaps, when the rule-wise epistemically vicious act (like wishful thinking) happens to get the truth, our intuitions misfire on the value of the result.

Alvin Goldman makes a comparable response to mine in his half of a forthcoming paper co-authored with Erik Olsson. Like me, he suggests that we may be confusing the instrumental value of a reliable process with greater final value of the product, and like me he alludes to money as an analogous case. The difference, though, is that Goldman suggests that knowledge actually does get more value, through a process of "value autonomization". This is a psychological thesis about how we come to *attribute* independent value to otherwise instrumentally valuable things, combined with the view that such intuitive attribution is a guide to what's *actually* of value. In the cases at hand this latter claim is implausible, though; we surely don't think such intuitions are a guide to what's actually of final value when it comes to money, for example. And though I'm sympathetic to letting intuitions be a defeasible starting place, they generally no longer have the same evidential force when we have an undermining story about their source—just as stubborn intuitions about which line is longer in the Müller-Lyer illusion do not remain evidential after learning the nature of the illusion.

Unlike Goldman, the pure utilitarian can simply let go of any attempt to make sense of knowledge as of more final value, and use the typical utilitarian story—one to which they're already committed—to do so. Knowledge is not more valuable than mere true belief, just as earnings are not more valuable than windfalls. In both cases the former seem more valuable because it makes sense to seek them, while seeking the latter is at best unwise and at worst incoherent.

### 3 Justification and epistemic responsibility

The theory of value behind an epistemology also has implications for the nature of epistemic justification. Two central characteristics of epistemic justification are that a) it's a positive evaluation of a belief, and that b) this positive evaluation can obtain even when the resulting belief is false. One typical and natural way to capture both these aspects is to say that justified beliefs are those that have been *responsibly* formed. Justification is then a positive evaluation because it represents epistemic responsibility on the part of the thinker, and yet we know that despite such responsibility, an uncooperative environment can still cause a thinker to fail to get knowledge.

The typical understanding of “responsibility” here involves deliberative cognition of some kind, but this construal has at least two serious problems. First, responsibility may then seem unnecessary for justification, since many are willing to make justification-like positive evaluations of beliefs even when they clearly *weren't* under the thinker's deliberate control. It is hard to say the thinker is “responsible” for beliefs that result from straightforward perceptual mechanisms, for example, but such beliefs (even when false) seem importantly well-formed in a way that at least looks like typical epistemic justification. In effect, taking justification as epistemic responsibility seems to beg the question against the justification externalists, who (roughly speaking) deny that justification supervenes on deliberative cognition. Second, such a notion of epistemic responsibility seems to require a robust doxastic voluntarism. If there is no such thing, as many believe, then epistemic responsibility (and thus justification, here) is not even *possible*.

These problems should sound like special cases of more general problems in the intersection of action theory and normative ethics. Is responsible choice in intention what's essential to positive evaluation? Is there such a thing as a choice for which we can be (truly) responsible in the first place? The Kantian deontologist must answer both of these positively, but the utilitarian need not. This is especially good for the epistemic utilitarian, since the conceptual coherence of free will faces tougher challenges than usual at the doxastic level.<sup>14</sup> Traditional libertarian views

---

<sup>14</sup>The *locus classicus* on doxastic voluntarism as an objection to deontic epistemology is Alston (1988). For a

will ascribe some separate form of causation to agents from the natural causation to which we're accustomed, and as applied in epistemology this mysterious causation would foreclose on any scientific account of how we form beliefs. It is far too early in the science side of the game for such foreclosures, and indeed given how far the science already has come, such foreclosures would look entirely arbitrary.<sup>15</sup> It's possible a compatibilist version of doxastic voluntarism would be sufficient for Kantian epistemology, but compatibilism, too, has worse than usual problems at the cognitive level. If we explain epistemic agency in terms of causal processes, our inclination to epistemic blame is even more likely to shrivel than in the ordinary action cases. A compatibilist approach to epistemic agency will require a close look at the causal mechanisms at the sub-propositional level that underwrite the formation of belief. These are not likely to be susceptible to, say, a Frankfurt-type account of freedom, since it will make little sense in most cases to speak of the agents' desires (of any order) with respect to such fine details of cognition.<sup>16</sup>

Again, at the level of normative ethics the utilitarian has a typical strategy for handling these problems—one that can be carried straightforwardly into the epistemic realm. Since utilitarian value does not depend on formation of free intentions, the utilitarian can consistently deny the existence of free will. The utilitarian also must deny the existence of what we might call “thick” moral responsibility—the moral responsibility of the retributivist, who claims that it can be just to punish even when net consequences of such punishment are bad.<sup>17</sup> To the utilitarian, punishment and reward are only justified, as anything, by its impact on welfare. The utilitarian thus has what we might call a “thin” notion of moral responsibility. Roughly speaking, for a utilitarian to say someone is *responsible* for an outcome is just to say that person houses the most appropriate place to apply change in order to bring about more future utility with respect to that outcome. When

---

collection on such issues, see Steup (2001).

<sup>15</sup>The naturalistic libertarian picture in Balaguer (2004), on the other hand, has problems explaining how we can be responsible for genuinely random occurrences in the brain.

<sup>16</sup>Frankfurt (1971).

<sup>17</sup>This is, I take it, the same thing as “desert”, or what Galen Strawson calls “true” moral responsibility—“responsibility of such a kind that, if we have it, then it *makes sense*, at least, to suppose that it could be just to punish some of us with . . . hell and reward others with . . . heaven” (Strawson, 1994, p. 216).

Squeaky is responsible for murder, this means that efforts to redress current and future such harms are best concentrated on her. Squeaky is not morally responsible in the thick sense for the act; she does not deserve punishment for its own sake. Nonetheless it can be just to subject her to isolation and rehabilitation for the sake of future consequences. To the extent we discover that Charlie *caused* Squeaky to murder, that is the extent to which we discover Charlie is the more effective place to apply such change, and thus in the utilitarian sense we attribute him more of the moral responsibility for the crime.

Because people are such complex creatures, they often appear as black boxes when it comes to their own causes, and thus it is usually at the level of people that we can most usefully place such moral responsibility. (When the cause is transparent, as in hypnosis cases, we then invariably place the responsibility further back.) Besides, people are literally more response-able than most other potential candidates for change. For these two reasons people are usually the place to look for assigning utilitarian responsibility. Similarly, people are often the most natural subject of assigning *epistemic* responsibility. That is, though sources of risk in belief are ubiquitous and varied, often the most salient way to *control* for belief risk is at the level of our own belief formation. It is so important that it just might for a long time seem to be the only salient place. This serves as the utilitarian explanation for the source of internalist intuitions—they are from judgements about epistemic risk-management. And just as it was revolutionary in criminal justice to suggest that sometimes it is most useful to blame a crime on the hypnotist, or the abusive parents, or the corrupting peers, or the hormonal imbalance, or the economy, or the government, or some other factor “external” to the criminal, so too was it revolutionary in epistemology to suggest that sometimes it is most useful to place epistemic blame or praise on factors external to the cognitive awareness of the thinker, as externalism suggests.

The utilitarian can thus take on the picture of epistemic justification as epistemic responsibility, with the understanding that for a thinker to be epistemically responsible is just to have controlled for risk in belief formation (with respect to epistemic goals) as best as might be expected—in other

words, so that no effort to change the thinker for better future (epistemic) results is appropriate. When we see what went wrong in the Gettier cases, for example, we do not feel the thinker needs to amend her epistemic ways. This utilitarian version of responsibility still captures the two key principles of justification, since a) the positive evaluation of epistemic justification comes from the high expected utility of good risk management techniques, and b) many risks cannot be controlled for, so that a belief responsible in this sense can still be false.

Here is another advantage of this account: unlike views committed to doxastic voluntarism, the utilitarian picture of responsibility does not fade as we look at the fine-grained cognitive details. In other words, the utilitarian version of justification is naturalistic in at least the sense that it plays nicely with cognitive science. Epistemic agents are more like entire investment banks than individual financiers; there is not one little epistemic homunculus sitting in the brain making all the epistemic investments, but instead an intricate organizational chart of departments and subdivisions, each node and sub-node of which performing specific duties aimed ultimately toward the goal of forming good beliefs. And when a major investment goes south at a financial institution, it is wise to dole out the responsibility at a fine-grained level of functional organization, rather than vaguely gesture at a whole department. As we learn more about the organizational chart of the human brain, we can similarly apportion epistemic responsibility with more accuracy.

## 4 Conclusion

The question of knowledge's value and the nature of epistemic responsibility are just two examples of how a background value theory can inform one's epistemology. My hunch is that utilitarian epistemology has other good fruit in its branches. For example, Pritchard (2005) explores implications of the simple hypothesis that *luck* is what makes the essential difference between knowledge and mere true belief. In a later paper he notes an important open question for his picture:

... I think it is crucial to say more about just why having non-lucky true beliefs is

so important to knowledge, and this extension of the project will inevitably lead into issues about epistemic value which are currently at the forefront of discussion.<sup>18</sup>

The utilitarian about knowledge seems to have a simple answer to this question—namely, that to rely on luck is to act on low expected utility, and this is a grave utilitarian sin. We seek what excludes luck in epistemology for the same reasons we seek what excludes luck in finance: something of high instrumental value is at stake.

I think the utilitarian approach also has potential for debates over skepticism. On this account to ask whether we ever have knowledge is like asking whether we ever really have earnings. Aren't all profits *really* just windfalls, given the countless ways even the apparently safest investment could have gone wrong? The utilitarian is insouciant about such questions; whether we call them windfalls or earnings, mere true belief or knowledge, what's important is the expected utility. What determines the threshold of a "safe enough" epistemic investment is a matter removed from the question of how to make it as safe as possible.<sup>19</sup> The utilitarian epistemologist is interested in the latter; like the financier, she will simply seek to act on what has the highest expected utility available, whether or not it counts as "safe enough".

If utilitarian epistemology can cash in on just some of the promise in this prospectus, then it will merit the commitment of further philosophical resources.

## References

Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2:257–299.

---

<sup>18</sup>Pritchard (2007a) p. 293.

<sup>19</sup>Here I do not mean the narrow, particularly epistemological sense of "safe" (according to which, roughly, a true belief is safe if it couldn't easily go wrong:  $Bp \square \rightarrow p$ ). I don't mean it as widely as "low risk", either. It means something like "high expected utility, not forgetting to account for reasonable risk aversion in diminishing marginal returns of utility on the instrumental good in question." Thus a 1 in 10 chance at a profit of \$10,000 might be a "safe" bet, though risky, while a 1 in a million chance at a \$1 billion profit is not safe, though of equal expected value.

- Balaguer, M. (2004). A coherent, naturalistic, and plausible formulation of libertarian free will. *Noûs*, 38(3):379–406.
- Chisholm, R. (1977). *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, NJ, second edition.
- DePaul, M. and Zagzebski, L., editors (2003). *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press, Oxford.
- Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 68(1):5–20.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23:121–123.
- Ginet, C. (1975). The general conditions of knowledge: Justification. In Alcott, L. M., editor, *Epistemology: The Big Questions*, pages 79–89. Blackwell Publishing, Malden, MA, 2004 edition.
- Goldman, A. I. (2000). Epistemic folkways and scientific epistemology. In Sosa, E. and Kim, J., editors, *Epistemology: An Anthology*, pages 438–444. Blackwell Publishing, Malden, MA.
- Greco, J. (2003). Knowledge as credit for true belief. In DePaul and Zagzebski (2003), pages 111–134.
- Harman, G. (1977). *The Nature of Morality*. Oxford University Press.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Plato (ca. 380 BCE). *Meno*. Public domain.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford University Press, Oxford.
- Pritchard, D. (2007a). Anti-luck epistemology. *Synthese*, 158(3):277–297.

- Pritchard, D. (2007b). The value of knowledge. In Zalta (1995). Online; last accessed April 2, 2008 from <http://plato.stanford.edu/entries/knowledge-value/>.
- Riggs, W. D. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64(1):79–96.
- Sosa, E. (1997). Reflective knowledge in the best circles. *Journal of Philosophy*, 94(8):410–430.
- Sosa, E. (2003). The place of truth in epistemology. In DePaul and Zagzebski (2003), pages 155–179.
- Steup, M., editor (2001). *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press, Oxford.
- Steup, M. (2006). The analysis of knowledge. In Zalta (1995). Online; last accessed July 3, 2008 from <http://plato.stanford.edu/entries/knowledge-analysis/>.
- Strawson, G. (1994). The impossibility of moral responsibility. In Watson, G., editor, *Free Will*, pages 212–228. Oxford University Press, second edition.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press, 2002 edition.
- Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy*, 34(1/2):12–28.
- Zalta, E. N., editor (1995). *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Stanford, CA.