

Trinity University

Digital Commons @ Trinity

Philosophy Faculty Research

Philosophy Department

2007

The Easy Argument

Steven Luper

Trinity University, sluper@trinity.edu

Follow this and additional works at: https://digitalcommons.trinity.edu/phil_faculty



Part of the [Philosophy Commons](#)

Repository Citation

Luper, S. (2007). The easy argument. *Acta Analytica*, 22(4), 321-331. doi:10.1007/s12136-007-0014-9

This Post-Print is brought to you for free and open access by the Philosophy Department at Digital Commons @ Trinity. It has been accepted for inclusion in Philosophy Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

THE EASY ARGUMENT

To say that knowledge is closed under entailment is to say that the following principle (perhaps with qualifications) is correct:

K: If, while knowing p , subject S believes q because S knows that p entails q , then S knows q .

This principle that knowledge is closed under entailment, K, has been challenged on the basis of cases like the following.

Table Case: Ted is in an ordinary house in good viewing conditions and believes *red*, his table is red, entirely because he sees his table and its color; he also believes *not-white*, it is false that his table is white and illuminated by a red light, because *not-white* is entailed by *red* (Stewart Cohen 2002).

Car Case: Sam has parked his car in typical (un-Gettierized) circumstances and believes *car*, his car is parked outside, because he just left it there; he believes *not-(not-car & dreaming)*, he is not merely dreaming his car is parked outside, because it follows from *car* (Gilbert Harman and Brett Sherman (2004).

Given the strength of his epistemic position vis-à-vis *red*, Ted seems to know *red*; similarly, Sam's epistemic position seems strong vis-à-vis *car*, so strong, in fact, that he appears to know *car*. But it may seem too *easy* for Ted to know *not-white* by deducing it from *red*, which he believes via perception; and too easy for Sam to know *not-(not-car & dreaming)* by deducing it from *car*, given his epistemic position vis-à-vis *car*. And these impressions are at odds with K, despite K's own obvious intuitive appeal. These examples and the like (which we may call *hard cases*) illustrate an interesting problem, namely, the following three claims clash but each seems plausible:

1. Ted's epistemic position is strong enough for him to know *red*.
2. Ted cannot know *not-white* on the basis of *red*.
3. The epistemic closure principle, suitably restricted, is true.

Other examples (discussed later) illustrate how intuition can suggest that our epistemic position is strong enough for us to know things that fail to position us to know other things for which the former provide powerful inductive support. Stewart Cohen (2002) has called this three-way clash of intuitions the problem of easy knowledge.

A skeptical response to the problem would be to accept 2 and 3 and reject 1.

Those who hope to avoid skepticism appear to have two options.

According to the *hard argument*, the best response is to reject K, and maintain that while Ted and Sam know *red* and *car*, they know neither *not-white* nor *not-(not-car & dreaming)*. A second response is to say that, despite appearances, Ted knows *not-white*, and Sam knows *not-(not-car & dreaming)*. Here we reject the assumption that in the hard cases knowledge of *not-white* and of *not-(not-car & dreaming)* comes too easily. Call this the *easy argument*. But there may be a third alternative. Perhaps we can eliminate the possibility that a belief can too easily be known on the basis of another by tightening our requirements for knowing the latter slightly, not enough to put ourselves in danger of substantial skeptical consequences, and without abandoning K. On this approach, we say the reason why it is too easy for Ted to know *not-white* by deducing it from *red* is that Ted does not know *red* to begin with, and similarly for Sam. We say that Ted's epistemic position vis-a-vis *red* does not suffice for knowledge, since he cannot know *red* merely by seeing that his table appears red, but he can easily improve his

epistemic position enough to know *red*. Thus what is needed is an account of knowledge that is strong enough to nip the possibility of easy knowledge in the bud, but not so strong as to prevent people from knowing things using their senses. Admittedly, this concedes ground to the skeptic, but perhaps the price is worth paying if we can avoid easy knowledge and retain K. I will call this approach the *reverse argument*.

In this essay I take on two tasks. In Part 1 I put aside the hard argument and criticize a recent version of the reverse argument. I claim that the reverse approach to the problem of easy knowledge leads back to skepticism after all. In Part 2 I criticize one version of the hard argument. My criticisms help support the easy argument, in that they chip away at its alternatives. However, all three arguments have awkward consequences. My thought is that it is easiest to live with the awkward consequences of the easy argument.

The Reverse Strategy

Assuming that knowledge can be analyzed in the way the reverse theorist expects (i.e., we can find an account that nips ‘easy’ knowledge in the bud without denying K and without substantial skeptical consequences), it will be possible to criticize any account that permits instances of easy knowledge, such as Ted’s knowing *not-white* in the Table Case. But a successful reverse argument against an analysis must do more than show that the analysis tolerates easy knowledge. For it is possible that *no* plausible account rules lives up to the expectations of the reverse theorist, and it is idle to object to an account on the grounds that it fails to do what no plausible account can. I will argue that the reverse strategy I consider does not live up to the expectations of the reverse strategists themselves. I will then suggest (but not demonstrate) that the failure was inevitable since any account that is strong enough to avoid easy knowledge is so strong as to have implausible skeptical consequences. The upshot is clear: its compatibility with easy knowledge is not grounds to reject an account.

I will consider a version of the reverse argument deployed by Richard Fumerton (1995) and Jonathan Vogel (2000) against various reliabilist accounts of knowledge. They object to reliabilism because it permits a pattern of reasoning Vogel calls ‘bootstrapping,’ and bootstrapping generates knowledge too easily. Vogel (p. 614) offers the following example. Roxanne “believes implicitly what her gas gauge says, without knowing that the gauge is reliable. . . . When the gauge reads ‘F’, she believes that, on this occasion, the tank is full. She also believes that, on this occasion, the gauge reads ‘F’.” Combining these, she believes that on this occasion the gauge reads ‘F’ and ‘F’ is true. This last proposition entails that, on this occasion, the gauge’s reading is accurate. Roxanne repeats her inference pattern again and again, and concludes, by induction, that the gauge is reliable. Vogel claims that at each step reliabilism implies that Roxanne knows that her beliefs are true. It implies that she can know her gauge reads accurately because a reliable process so indicates, namely the gauge reading itself. And, assuming induction is reliable, it implies that she can now put several such beliefs together so as to know her gauge is reliable. Since Roxanne’s bootstrapping is objectionable, yet permitted by reliabilism, we should reject reliabilism.

Vogel’s example is flawed. The reliability of a gauge is not reliably indicated merely by accurate readings on a number of occasions, no matter how large the number. Nor are such readings the basis for a saliently strong inductive inference—one capable of

generating knowledge. Maybe my gauge is stuck on ‘empty,’ and my car has been up on blocks, with an empty tank, for years, during which time I check the gauge twice daily, and each time the gauge reads ‘empty.’ My sample of readings is simply not representative. Only its repeated accurate readings in saliently diverse sorts of circumstances could reliably indicate its reliability and form the basis for an inductive inference that positions us to know the gauge is reliable.

To help Vogel out, we can reconstruct his example. The reconstruction begins like the original: Roxanne believes p because her gauge, which is reliable, indicates p . She gathers many similar beliefs, attained because of the gauge’s readings. She adds the premise that the readings were taken in saliently varied circumstances. She then infers, via induction, that her gauge is reliable.

Let us assume, for the time being, that reasoning involving bootstrapping is flawed, and that reliabilism tolerates bootstrapped knowledge. We have said that this is a strike against reliabilism only if there is a plausible way to rule out bootstrapped knowledge. Is there a way?

Vogel thinks there is. He wants to revive the traditional view that knowledge entails justification: Roxanne can know her gas level is such and such because her gauge says so only if she is justified in believing her gauge reliably indicates her gas level (622). What she needs is an “independent reason to believe that the position of the needle on the gauge is reliably correlated with how much gas is in the tank.” Since she must have such a reason at her disposal at the outset, she cannot bootstrap her way to knowledge. Consider some reservations about Vogel’s suggestion.

As Cohen notes, ‘easy knowledge’ is not limited to cases of bootstrapping. If it is too easy for Roxanne to know her gauge is reliable, it is also too ‘easy’ to know *not-white*, a table is not white with red light shining on it, by deducing it from *red*, the table is red, where the latter is believed spontaneously via perception, a reliable process. Assuming that knowledge in the Table Case qualifies as ‘easy,’ then Vogel needs a way to preclude it.

Vogel’s view would be this: what has gone wrong is that Ted does not have a justified belief that his knowledge source, which is his color vision, is reliable, hence he does not know it is reliable. Let’s adjust the example accordingly. Assume that Ted believes his vision is generally reliable, and that his belief is justified. Unfortunately, even under these circumstances, it seems that Ted too easily knows *not-white*. The adjusted example is not made unproblematic by Ted’s justified belief in the reliability of his vision.

Why does Ted’s knowledge still seem too “easy”? Because Ted’s being in a position to know *red* depends on the truth of *not-white*. Hence Ted’s knowing *not-white* upon deducing it from *red* seems suspiciously circular, and this appearance is not eliminated by the assumption that Ted has a justified belief that his color vision is generally reliable.

The point can be made clearer if we distinguish between two senses in which color vision might be reliable. Even though color vision is *generally* reliable, there are circumstances in which it is useless, and Ted’s knowing his table is red depends on his not being in such circumstances. For example, color vision does not work well in nonstandard lighting conditions, such as when a white table is illuminated by red light. There is a type of reliability it lacks in nonstandard lighting conditions, and only when it

has this type of reliability will it produce knowledge. For convenience, I will say it lacks *specific reliability*. Refinements aside, a source is generally reliable when the beliefs it endorses would be true if it were used in a wide variety of actual circumstances, while a source is specifically reliable when the beliefs it endorses would be true if it were used specifically in circumstances like those at hand. Specific reliability is necessary for knowledge (Luper 1987b). That is the upshot of Gettier's paper. The general reliability of Ted's vision did not depend on (or support) *not-white*. However, the specific reliability of Ted's vision *did* depend on the fact that *not-white*. Hence the assumption that vision is generally reliable does not remove the appearance of circularity involved in Ted's knowing *not-white* on the basis of *red*.

(Conceivably, Vogel might respond by claiming that, to know *red*, Ted needs to know, hence justifiably believe, that his color vision is *specifically* reliable. However, as I will suggest in below, down this road lies skepticism.)

Vogel's reverse strategy has not succeeded; he has not provided us with a plausible way to nip easy knowledge in the bud. In part this failure is due to the fact that he has underestimated the problem of easy knowledge. In what follows I will attempt to characterize its essential feature. If I am correct, we cannot rule out the possibility of easy knowledge without either abandoning closure or adopting an analysis with unacceptable skeptical consequences. Assuming that neither alternative is acceptable, we can conclude that sometimes knowledge just *is* easy.

So why does it seem counterintuitive to say that Ted and Roxanne know that the things they believe are true? The best explanation, I suggest, turns on the fact that these cases seem to involve reasoning from a proposition to something that *grounds* that very proposition, in this sense: *g grounds p for person S* just in case *g's* truth is instrumental to *S's* knowing *p*. Let us say that such reasoning is *pseudocircular* (Luper 2005 and 2006). It seems counterintuitive to say that knowledge can depend essentially on pseudocircular reasoning, as it would in the case of Ted and Roxanne, and therefore counterintuitive to attribute knowledge to Ted and Roxanne. (To allow for the possibility of believing something through multiple sources, we should put the explanation this way: knowing a belief is true requires having at least one source that does not involve pseudocircular reasoning, yet Ted's and Roxanne's beliefs have no such source. In the interest of simplicity, I will not pursue this alternative explanation.)

My explanation makes reference to truths that are "instrumental to" our knowing things. I choose this admittedly vague terminology deliberately, so that my explanation will not presuppose the truth of any particular theory of knowledge. Different theorists will have different views about when it is that a proposition's truth is instrumental to one's knowing things. By way of illustration, consider the following points about the Table Case, which, I think, are fairly uncontroversial. Ted's knowledge source is roughly his visual process. By this process Ted knows things only if under his circumstances vision is sufficiently reliable. Its being sufficiently reliable depends on the truth of various propositions; each such proposition is instrumental towards Ted's knowing things through his source. An example is the proposition that it is false that Ted's table is white and illuminated by a red light. Proposition *g* grounds our knowing *p* when *p's* source's requisite reliability hinges on *g's* truth. Next consider propositions that defeat reasoning that is essential to someone's believing *p* (without justifying a false belief thereby): the

negation of any such defeater is instrumental towards her knowing p on the basis of that reasoning.

I know of no clearer general account of the propositions that ground knowledge. The possible accounts that come to mind seem flawed. For example, suppose we say that g grounds p for person S just in case:

(S knows p) entails g .

But I know *fishes*, fishes live in water, and my knowing *fishes* entails each of the many things which *fishes* itself entails, such as that either fishes live in water or the moon hit my eye like a pizza pie, yet few—certainly not all—of these propositions play a role in my knowing *fishes*.

A better account says that g grounds p for person S just in case:

(S knows p) entails (or perhaps materially implies) g but p does not entail g .

Yet this account eliminates propositions which appear to be implicated in the hard cases. For example, Sam's knowing *not-(not-car & dreaming)* on the basis of *car* is considered too easy by many theorists, such as Harman, even though the latter entails the former.

Or should we say that *not-dreaming*, and not *not-(not-car & dreaming)*, grounds Sam's knowledge that *car* holds? We might say that the former (like my brain is not having manufactured experiences) is the negation of a *core* skeptical hypothesis, and the latter (like I'm not a detached brain *on the far planet Crouton* having manufactured experiences) is the negation of a trivial consequence of that hypothesis. Perhaps Sam's knowing that *car* holds is grounded by the falsity of the core skeptical hypothesis, but not by *not-(not-car & dreaming)*. We have already said that a proposition may play a role in our knowing something even if some of its consequences do not. Suppose we say that while knowing things may be grounded by the falsity of core skeptical hypotheses, it is not grounded by the falsity of their trivial consequences. If something like this core thesis holds, we can dismiss the claim that Sam too easily knows *not-(not-car & dreaming)* as an illusion resulting from a failure to see that this proposition is not really instrumental to his knowing *car*. Knowing *not-(not-car & dreaming)* wholly on the basis of one's knowledge that *car* holds is no more problematic than knowing I have at least one hand wholly on the basis of my knowledge that I have two. We can also say that the Table and Car Cases are no threat to K. Given K, Ted better be in a position to know *not-white* if he knows *red*, and Sam better be set to know *not-(not-car & dreaming)* if he knows *car*, but that is no problem if the former do not ground the latter. Ted may be in no position to know that his table is not being illuminated by red light even though he knows *red*, and Sam may be unable to know that he is not dreaming although he knows *car*, since (among other things) the former ground the latter, but that is entirely consistent with K (contrast K with Moore's principle, discussed in Luper 2007). I expect that most theorists will not accept the core thesis, and I will not rely on it in what follows.

The hard cases will be problematic to those who think that a proposition cannot be known on the basis of something it grounds (and who reject the core thesis). In Vogel's example, it is bootstrapping that is pseudocircular. A belief's truth can be known only if its source is generally reliable. Roxanne's gauge's readings position her to know things only if its readings are reliable. So if these things (which her gauge tells her) are her reasons for believing that her gauge is reliable, her reasoning is pseudocircular.

I suggest that if we accept K, and reject skepticism and the core thesis, we will have to tolerate pseudocircularity (compare Van Cleve 2003). To rule out

pseudocircularity compatibly with K, we will have to accept something like the following *Independence Condition*:

If S knows *k*, and proposition *g* is instrumental to S's knowing *k*, then S knows *g*, and *k* is not instrumental to S's knowing *g*.

Equivalently:

If proposition *g* grounds *k* for S, then S knows *g*, and *k* does not ground *g* for S. The Independence Condition blocks pseudocircular sources of knowledge: given the Independence Condition, nothing that grounds a bit of knowledge may be known on the basis of that bit of knowledge. The Independence Condition is also consistent with closure: each consequence of a bit of knowledge *k* must be known independently if it grounds *k*, but may be known on the basis of *k* if it does not ground *k*.

Instead of the Independence Condition, why not adopt a weaker principle that allows us to know a proposition *k* without independently knowing the truth of something *g* that grounds *k* so long as *k* does not entail *g*? The following principle is weaker in precisely this way:

If proposition *g* grounds *k* for S, and *k* entails *g*, then S knows *g*, and *k* does not ground *g* for S.

This, the *Consequence Independence Condition*, like its predecessor, precludes Ted's knowing it is false that his table is white and illuminated by red light because it follows from the fact that his table is red. Yet the Consequence Independence Condition allows Ted to know his table is red without independently knowing that it is not illuminated by red light, even though the latter grounds the former, since the table's being red does not entail that it is not illuminated by red light. However, it is difficult to see how the exceptions allowed by the Consequence Independence Condition would be motivated: if knowing that his table is red requires his independently knowing that it is not white and illuminated by red light, why shouldn't it also require his independently knowing that it is not illuminated by red light? (As noted earlier, if we were to say that only one of these two grounds Ted's knowing *red*, it would be more plausible to pick the table's not being illuminated by red light, rather than the falsity of its being white and illuminated by red light.)

I have said that the Independence Condition allows us to ban pseudocircularity while retaining closure. Unfortunately, however, if we rely on the Independence Condition to reconcile closure with the ban on pseudocircularity, we must pay a price: namely, skepticism concerning ordinary cases of knowledge.

Here's why: The Independence Condition requires that, to know his table is red, Ted must know that his vision is specifically reliable, and that none of the circumstances in which it is not specifically reliable obtain. It also requires that he knows these things in a way that is independent of any knowledge which his vision gives him with their help, such as his knowledge that his table is red. But that means Ted does not know his table is red, since he *didn't* try to eliminate the many possibilities that would undermine the specific reliability of his vision. For example, he did not attempt to establish that he is not having visual hallucinations involving tables. (In the Table Case Ted did conclude that it is false that his table is white and illuminated by a red light, but he did not try to establish this independently of his knowledge that his table is red.) Now consider Sam. To establish *not-(not-car & dreaming)*, Sam must establish either *car* or *not-dreaming*. To do so without relying on *car*, he will need to establish *not-dreaming*. But he didn't

even attempt this. So he does not know *car*. Yet on anyone's list of ordinary cases of knowledge, Ted's belief *red* and Sam's belief *car* would appear. And, like Ted and Sam, most of us, most of the time, do not really know ordinary empirical truths, since we don't try to rule out possibilities such as *not-dreaming*.

That most of us most of the time do not know commonsense empirical truths is an extremely counterintuitive skeptical consequence. Of course, things could be worse. It could turn out that we *cannot* know such truths. And in fact the Independence Condition might well make it impossible for us to know ordinary empirical truths. In effect, the Independence Condition demands, of any knowledge source, that it be checked out, in the sense that we must come to know that it is specifically reliable, and that the truths on which that reliability depends hold. If all putative knowledge sources have to be checked out in order for them to give us knowledge, where will we get the knowledge to do the checking? Here we face the standard skeptical trilemma. Our efforts to check our sources will begin with assumptions whose truth we do not know, or regress indefinitely, or they will involve some sort of circularity which the Independence Condition rejects. Consider that all five senses will lack specific reliability if certain skeptical hypotheses hold, such as our suffering a complex set of hallucinations affecting all of our senses. Given the Independence Condition we cannot use visual knowledge to verify that these hallucinations are not undermining our visual sense. Nor may we use tactile knowledge to confirm that the hallucinations are not undermining our tactile sense. Can we use visual knowledge to verify that they are not undermining our *tactile* sense, and tactile knowledge to verify that they are not undermining our visual sense? Apparently not; to have visual knowledge we must not be suffering the hallucinations; if we use our visual knowledge to verify something that in turn confirms that we are not suffering the hallucinations, we violate the Independence Condition.

The Hard Argument and Lotterylike Propositions

Reverse arguments, such as the one rejected in the previous section, are relatively recent additions to the epistemological literature. Hard arguments came first. One version of the hard argument is very well known (Robert Nozick 1981; Fred Dretske 1970, 2003, 2005): Let us say that a proposition is *elusive* if and only if our experiences would remain the same if the proposition were false. For example, *not-white* is true, and Ted has certain experiences which he would still have if *white* were true. According to the argument from elusiveness, we fail to know of elusive propositions that they are true even if we believe them because we see that they are entailed by things we know, so we should reject K. By now the argument from elusiveness is well criticized. I will not discuss it further. Instead, I will consider a version of the hard argument that attempts to use lottery propositions and lotterylike propositions against the principle of closure.

The paradigm case of a *lottery* proposition is *not-win*, the ticket in my hand—one of the ten million issued in the state lottery that will end tonight—is not the winner. What is distinctive about these propositions is that, normally, they are supportable only on the grounds that they are highly likely. For example, to support my claim that my ticket will lose, normally I would cite the fact that the probability is very high, albeit less than 1. As Jonathan Vogel (1990) and other theorists (see especially Hawthorne 2006) have noted, some propositions that do not actually involve lotteries still resemble lottery propositions in that they can be assigned a probability that is less than 1. Let us say that

these propositions are *lotteryesque*. For example, *not-stolen*, my automobile has not been stolen and taken south of the border, seems *lotteryesque* given the statistics concerning stolen vehicles in the U.S., relative to which the probability of *not-stolen* is less than 1, even if very high.

Lottery propositions cannot be known solely on the grounds that their truth is highly likely (Harman 1968). To insist that they can be known on this basis raises the specter of Kyburg's (1961) lottery paradox. Consider, too, that it is unacceptable to say both 'I know *p*' and '*p* might be false,' in the epistemic sense of 'might.' Yet any lottery proposition might be false in the epistemic sense.

As several theorists have noticed, we can avoid paradox and explain why we normally fail to know lottery propositions if we say that knowing *p* requires believing *p* because of something that *establishes p's* truth. This view of knowledge has received different but closely related formulations: Dretske (1971) said knowledge requires having a conclusive reason for thinking that what we believe is true; David Armstrong (1973, p. 187) said "knowledge requires a belief-state which *ensures* truth;" and Sherman and Harman (2004, p. 492) say "one knows only if one believes as one does because of something that settles the truth of that belief." Proponents of the safe indication account of knowledge (e.g., Luper 1984, 2003a; Sosa 2000) will also say we know things only if we believe as we do on grounds that establish truth. On this account we know *p* only if we believe *p* on the basis of an event or state of affairs R that safely indicates *p's* truth, where R safely indicates *p's* truth only if the following subjunctive conditional is true:

p would hold if R held.

On each of these approaches, we fail to know things, including lottery propositions, when our sole basis for believing them is their high likelihood.

It is presumably because normally we simply do not know lottery propositions that some theorists consider them 'hard.' Their hardness is not much of a threat to K, however, since it is not obvious that there are mundane knowledge claims that entail genuine lottery propositions. Consider *not-buy*, I will not buy a 10 million dollar villa in the French Riviera tomorrow, since I lack the means, and the conditional, if *win* then *buy*, i.e., tomorrow I will buy the villa if I win the state lottery tonight. If *not-buy* and if *win* then *buy* are among the things I know, K is under pressure, since these entail *not-win*, so that, given K, I can easily know *not-win*. More precisely, what is under pressure here is not K but rather the following stronger principle:

GK: If, while knowing various propositions, S believes *p* because S knows that they entail *p*, then S knows *p*.

But the proponent of GK is well positioned to argue that I do not know *not-buy*. One reason I fail to know it is precisely that its truth depends, in part, on whether I win the lottery, and I do not know I will not (compare Harman 1986, p. 71).

However, this strategy fails as applied to the wider group of *lotteryesque* propositions. It is relatively uncontroversial that I know *ford*, my 1969 Ford 100 is parked in my garage downstairs. But *ford* entails the *lotteryesque* proposition *not-stolen*, so that, given K, the latter is easily known. And friends of K cannot plausibly respond by denying that I know *ford*. Too many of the propositions that we quite clearly know entail *lotteryesque* propositions.

A better strategy is to emphasize that genuine lottery propositions are normally supportable only on the grounds that their truth is highly likely whereas *lotteryesque*

propositions may be supportable on grounds that establish their truth. While lottery-esque propositions *can* be based on probabilistic grounds, their truth cannot be known on such grounds. But they can also be believed because of something that establishes their truth, and hence they can be known. My belief *not-stolen* is not based on crime statistics; if it were, I would not know that it is true, since on this basis my belief is at best highly likely. Featuring prominently among my grounds is my observation *O*: I only just parked my Ford downstairs. It would not be true that *O* establishes that *not-stolen* holds in Gettierized circumstances; for example, *O* would not do the trick if there were car thieves at work in my neighborhood, if I had a son with his own pair of keys who, unbeknownst to me, is about to drive off in my car, and so forth. But under common circumstances, *O* establishes that *not-stolen* and that *ford* holds.

We may draw a similar conclusion about genuine lottery propositions. They cannot be known to be true if believed solely because they are highly likely. But they can be known in unusual circumstances. To know *not-win*, I would have to know my ticket is counterfeit, or that the lottery is rigged against me, or the like. When S believes *p* upon seeing (knowing) it is entailed by something S knows, let us say that *p* is *knowledge secured*. Lottery-esque propositions are rarely knowledge secured, but when they are, their truth is known.

Acknowledgements

A brief version of this essay was presented at Bled Conference on Epistemology, May 28-June 2, 2007. I thank the participants, my colleague Curtis Brown, and anonymous referees for their comments.

References

- Armstrong, D., 1973, *Belief, Truth and Knowledge*, Cambridge: Cambridge University Press.
- Cohen, S., 2002, "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research* 65.2: 309-329.
- , 2005, "Why Basic Knowledge is Easy Knowledge," *Philosophy and Phenomenological Research* 75.2: 417-30.
- Dretske, F., 1970, "Epistemic Operators," *Journal of Philosophy* 67: 1007-1023.
- , 2003, "Skepticism: What Perception Teaches," in Luper 2003b.
- , 2005, "Is Knowledge Closed Under Known Entailment?" in Steup 2005.
- Fumerton, R., 1995, *Metaepistemology and Skepticism*, Lanham, MD: Rowman and Littlefield.
- Harman, G., 1968, "Knowledge, Inference and Explanation," *American Philosophical Quarterly*, 5: 164-73.
- , 1986, *Change In View*, Cambridge: MIT Press; partially reprinted in S. Luper, *Essential Knowledge*, New York: Pearson Longman, 2004.
- Harman and Sherman, 2004, "Knowledge, Assumptions, Lotteries," *Philosophical Issues* 14: 492-500.
- Hawthorne, J., 2006, *Knowledge and Lotteries*, Oxford: Oxford University Press.

- Kyburg, H. (1961) 'Conjunctivitis', *Probability and the Logic of Rational Belief*, Middletown, Conn.: Wesleyan University Press.
- Luper, S., 1984, "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism," *Australasian Journal of Philosophy* 62: 26-50.
- (ed.), 1987a, *The Possibility of Knowledge: Nozick and His Critics*, Totowa, NJ: Rowman and Littlefield.
- , 1987b, "The Causal Indicator Analysis of Knowledge," *Philosophy and Phenomenological Research* 47: 563-587.
- , 2003a, "Indiscernability Skepticism," in S. Luper 2003b.
- , 2003b, *The Sceptics*, S. Luper (ed.), Hampshire: Ashgate Publishing, Limited.
- , 2005, "Epistemic Closure Principle," *Stanford University Encyclopedia of Philosophy*, Edward N. Zalta, Editor, 2002.
- , 2006, "Dretske on Knowledge Closure," *Australasian Journal of Philosophy* 84.3 (2006) 379-394.
- , 2007, "Re-Reading G. E. Moore's 'Certainty,'" *Philosophical Papers* 36.1: 151-163.
- Nozick, R., 1981, *Philosophical Explanations*, Cambridge: Harvard University Press.
- Sosa, E., 2000, "Neither Contextualism Nor Skepticism," in Luper 2003b, 165-182.
- Steup, M. and Sosa, E. (eds.), 2005, *Contemporary Debates in Epistemology*, Blackwell, Malden, MA.
- Van Cleve, J., 2003, "Is Knowledge Easy—or Impossible? Externalism as the Only Alternative to Skepticism," in Luper 2003b, 45-61.
- Vogel, J., 1990, "Are There Counterexamples to the Closure Principle?" in *Doubting: Contemporary Perspectives on Skepticism*, M. Roth and G. Ross (eds.), Dordrecht: Kluwer Academic Publishers.
- , 2000, "Reliabilism Leveled," *Journal of Philosophy* 97: 602-623.

THE EASY ARGUMENT

Steven Luper
Philosophy Department
Trinity University

4-23-2006; 4-21-2007; submitted for publication 9-3-2007; revised 11-27-2007