

teorema

Vol. XXXIII/3, 2014, pp. 159-170

ISSN: 0210-1602

[BIBLID 0210-1602 (2014) 33:3; pp. 159-170]

Standing Beliefs, Skepticism, and Some Questions about Zalabardo's Probabilistic Approach

Adam Leite

I

Jose Zalabardo's *Scepticism and Reliable Belief* is bold and liberating in its forthright refusal to seek a single, uniform account of knowledge. It offers significantly different accounts of adventitious non-inferential knowledge, inferential knowledge, and knowledge constituted by standing beliefs, and no single unifying story ties together the requirements applying to the different cases. This strikes me as the expression of an admirable aspiration to take the phenomena as they are found, an aspiration likewise expressed in various methodological asides throughout the book. At the same time, however, I have reservations about a central aspect of Zalabardo's project. His proposals are articulated in terms of a highly developed mathematical apparatus for dealing with, e.g., a probabilistic measure of confirmation. But it remains quite unclear to me how this apparatus is to be applied to fairly straightforward cases, and as a consequence I have begun to wonder whether it captures anything in our actual epistemic lives at all.

I will begin with Zalabardo's account of knowledge through "standing beliefs" and its relation to skepticism. This will lead to some qualms, both substantive and methodological, about Zalabardo's new skeptical argument. Crucial here is his principle *PI*, which holds that "*S* can have inferential knowledge of *H* based on the evidence provided by *E* only if *S*'s belief in *E* confirms *H*" [p. 98]. As I will discuss, it is in attempting to apply this principle that I lose my grip both on how Zalabardo's probabilistic theory is to be applied and indeed on what his "probabilities" measure in the first place. I will close by returning to the issue of skepticism.

II

Zalabardo holds that standing beliefs – which he identifies as beliefs that we have an "innate predisposition to form ... that is largely independent

of input” [p. 137] – are knowledge whenever (and simply because) they are true [p. 138]. Among such beliefs, he includes the belief that one is not a brain-in-a-vat (BIV). Your knowledge that you are not a BIV is thus secure simply because you aren’t one. So much for the familiar skeptical worries.

Zalabardo claims that not much defense is needed for this account [p. 139]. This strikes me as a mistake. If different requirements are imposed regarding different types of knowledge, we naturally want some characterization of why different requirements are apt in the different cases. Why aren’t probabilistic constraints, which Zalabardo insists on in other cases, relevant in this case? The question would look less pressing if all the beliefs that we have an “innate predisposition to form ... that is largely independent of input” were true. But we know that they aren’t. We have an innate predisposition to believe, largely independently of input, in both the Gambler’s Fallacy and in the Euclidian structure of space. So our “innate predispositions” don’t always point in the right direction. Zalabardo introduces his notion of “standing beliefs” by suggesting that these “innate predispositions” are selected for by evolution [pp. 136-7]. Perhaps this suggestion is meant to help explain or reinforce their positive epistemic status. But innate belief-forming predispositions would not be selected for on the basis of truth, but rather for value to survival and reproduction, which is a different matter. Some account is consequently needed to explain why a belief’s being the result of an evolutionarily-selected “innate predisposition” is related to the belief’s truth in such a way that if the belief is true, it should count as knowledge. Moreover, it is an open empirical question whether every innate belief-forming predisposition is evolutionarily selected for at all. A particular belief-forming predisposition might rather be a spandrel of something else that is evolutionarily selected for. If that were so, then too we would wonder why the fact that the beliefs so formed also happen to be true should suffice for them to count as knowledge. Moreover, such phenomena suggest that Zalabardo’s way of thinking about standing beliefs is vulnerable to something much like “Gettierization”: a standing belief could be evolutionarily selected for because of evolutionary pressures having nothing to do with truth, and yet might, fortuitously enough, be true in the world as we find it. Is that really sufficient for knowledge? More discussion is needed.

Turning now to Zalabardo’s treatment of skepticism, I am not at all sure that the beliefs that we are not dreaming, not BIVs, and not being deceived by an Evil Demon *are* standing beliefs, given Zalabardo’s characterization. For one thing, I doubt that my belief that I am not dreaming right now is “largely independent of input”. I wouldn’t believe as I do if my experience were of a very different nature right now, and there are relevant features of my current experience to which I am appropriately responsive [Leite (2011a)]. On my view, in fact, we *can* detect the truth value of these anti-skeptical claims, at least insofar as we can have good evidence in their favor. I will return to this

point later. Right now, though, there is another point to stress. This is that I *do not* have an evolutionarily selected “innate predisposition” to believe that I am not a BIV. What I do have, if I have any innate predispositions in this territory, is an innate predisposition to believe that my sensory apparatus is hooked up to the world in more or less reliable ways, or that special cases aside, my senses are largely trustworthy. Of course, I also recognize that this wouldn’t be so if I were a BIV. So I believe that I’m not a BIV as well. This suggests that my knowledge that I am not a BIV must in fact be *inferential* on Zalabardo’s account. If that is right, however, then a problem looms. For now we have to apply Zalabardo’s crucial principle *PI* to this case, as to all inferential knowledge on Zalabardo’s account. As Zalabardo glosses it, this principle requires that you “be less likely to believe your evidence if the hypothesis were false than if it were true” [p. 98]. But would you be less likely to believe your senses to be reliable if you were a BIV than if you were not? Presumably not. So *PI* isn’t satisfied, and so you can’t know inferentially on the basis of your standing beliefs that you aren’t a BIV.

Maybe, however, the structure here isn’t really inferential. Maybe instead we have an innate predisposition to form something like *substitution instances* of a general schema, and each of them directly counts as standing knowledge so long as the proposition so-believed is true. However, this proposal points towards a possible response to Zalabardo’s own skeptical argument. A centerpiece of Zalabardo’s book, almost in fact its *raison d’être*, is his presentation of a new skeptical argument regarding our ability to know, of our true beliefs, that they are true. One possible rejoinder would be to suggest that we have an innate or at least standing predisposition to believe, of each of our beliefs, that it is true – i.e., that we have not a general standing belief to the effect that all of our beliefs are true, but rather a standing predisposition to form non-inferentially an appropriate instance of such a belief with regard to each of our beliefs. With this suggestion on the table, a possible account of our knowledge of the truth of our beliefs would be ready to hand. Just say that this belief-forming mechanism yields knowledge of the truth of one’s belief that *p* whenever one knows that *p* is true. On this proposal, this knowledge would arise from a mechanism that is *belief-dependent* — though not inferential — insofar as it gives rise to some beliefs given that one has certain other beliefs. But it is a mechanism that is truth-preserving, since whenever the input belief is true, the output belief will be true as well. I don’t see why someone with Zalabardo’s externalist leanings should think that it wouldn’t be knowledge-preserving as well. Of course, the mechanism will only be conditionally reliable. But that hardly seems a reason to deny that the output beliefs can constitute knowledge when the input beliefs do, especially for someone like Zalabardo who grants that conditionally reliable inferential processes can yield knowledge.

Maybe Zalabardo wouldn't like this account. But why not? What relevant condition on knowledge is violated? An important question this raises is how Zalabardo is thinking about "kinds" or "types" of knowledge and the requirements that apply to each. Zalabardo claims that my belief that I am not a BIV counts as knowledge simply in virtue of being true. If that isn't ad hoc, why should it be thought ad hoc to claim that I know that my belief that p is true simply because (1) I have an innate predisposition to form such a belief (given that I believe that p) and (2) I know that p ? What considerations license Zalabardo's view about our knowledge that we are not BIVs but rule out this proposed view about our knowledge that our beliefs are true? And even if the conditions Zalabardo insists upon for some other case would be violated here, why should that matter? After all, on Zalabardo's view different requirements apply to different sorts of knowledge: he already holds that standing knowledge violates conditions that apply to other sorts of cases, so why not extend the same liberality here?

III

I emphasize these questions because there is something about Zalabardo's methodology that puzzles me. His method – depicted with a very broad brush – is to seek a theory that yields, so far as possible, the results we pretheoretically deem correct, with minimal disruption to our commitments elsewhere and without imposing any broad overarching principles to unify our judgments about various sorts of cases. That approach is salutary, I think. But I worry that it can lead to a kind of partiality – a matter of picking and choosing – *given Zalabardo's aim of offering a surprising skeptical argument to the effect that we cannot know that any of our beliefs are true*. Zalabardo's overall skeptical argument amounts to this: "Look, I've set out conditions on three sorts of knowledge — inferential knowledge, adventitious non-inferential knowledge, and knowledge through standing beliefs – and our beliefs to the effect that our beliefs are true *do not meet the conditions contained in any of these three accounts, so they aren't knowledge*." But why should we think that Zalabardo's three "sorts" carve up the territory exhaustively? Why should we think that even if they do, there aren't sub-varieties involving different sorts of requirements? How are we to make a principled decision about these matters? A way must be found to do so, if we are to avoid the appearance of special pleading on behalf of Zalabardo's skeptical argument. For surely, given Zalabardo's overall methodology, the pre-theoretical judgment that we *do* know that some of our beliefs are true should carry a lot of weight. Given Zalabardo's methodology, it would seem that the interesting philosophical task would be to explain *how* we have this knowledge – what it

is that makes these beliefs have the status of knowledge on occasion – if they do not meet the conditions set out by the account so far.

Zalabardo attempts to respond to qualms on this score, commenting that “it might not be the job of the theory of knowledge to save” the judgment that we can know our beliefs are true, because “treating the conflict with epistemic optimism as a refutation of my account of knowledge begs the question against the anti-realist conception of the problem” [p. 165]. He goes on to claim that “the skeptical argument that my account of knowledge underwrites puts genuine pressure on realism” [Ibid.]. But this conceptualization of the dialectical situation is correct only if there is no plausible way of maintaining realist assumptions while rejecting Zalabardo’s skeptical argument. His argument amounts to the claim that because our beliefs about the truth of our beliefs do not meet the conditions he has set out for his three kinds of knowledge, those beliefs can’t constitute knowledge. But this argument can be rejected if we can sketch a way in which we know, on occasion, that our beliefs are true: there may be more things in heaven and earth than are dreamt of in this epistemology. If the additional account preserves realist assumptions, we won’t have objectionably begged the question against the anti-realist. We will have shown that even if Zalabardo’s skeptical argument presupposes realist assumptions, those assumptions need not be what is at fault.

IV

How, then, do I know that my beliefs are true? I don’t have a general answer to this question, and I am grateful to Zalabardo for raising it so forcefully. Earlier I offered one proposal broadly consonant with Zalabardo’s sensibilities. However, I’d also like to make a more fundamental point. It seems to me that Zalabardo’s approach to this question suffers from not attending to certain clear facts about how we proceed in everyday epistemic life. In fact, it seems to me that Zalabardo’s overall theory can’t make sense of a centrally relevant aspect of ordinary epistemic life.

What I have in mind is the familiar activity of *double-checking* to make sure that our beliefs are correct. Suppose that you and I are having dinner out together. We decide to leave an 18% tip. (Waitstaff are paid very poorly in the United States.) Calculating the tip using pencil and paper, we reach divergent results. We are each quite confident of our abilities and calculations, however, and despite the divergence we each believe that the amount we calculated is 18% of the total bill. After a little friendly teasing back and forth about our confidence in our abilities, we decide to check our calculations with a calculator. The result is precisely as I had said. “You see,” I announce triumphantly, “The answer is \$12.83. So I was right all along.”¹ Since I was perfectly well aware of what I believed throughout this process, “I was right

all along” here amounts to saying that my belief was correct all along. I had believed that the total was \$12.83, and I continued to believe it even as we double-checked, and I believe it now, and I have found my belief to be correct through double-checking. So it looks as if this is one perfectly ordinary way, even if it is not the only way, in which we can come to know that a belief is true. What exactly it involves is obscure, however. It seems to involve an inference, namely, from the facts (1) that p and (2) that I believe that p , to the conclusion that my belief is true. Similar cases with this form involve visual double-checking, asking an authority for confirming information, etc. In each case, I check whether p in order to confirm that my belief that p is true. When all goes well, I thereby know that my belief is true.

Zalabardo’s view can’t make any sense of this. Zalabardo argues that one cannot know inferentially (on the basis of the fact that p) that one’s belief that p is true, because the requirement *PI* on inferential knowledge would not be satisfied. *PI* would not be satisfied in this case, because – as he interprets it – it requires that in considering the relevant probabilities, we have to hold fixed that one believes p to be true. So in the case in which one’s belief that p is false (i.e., one’s hypothesis – that one’s belief that p is true – is incorrect), one still believes p to be true, and so one *ipso facto* is just as likely to believe one’s evidence (namely, that p is true) [pp. 155, 103-4]. So one cannot know, via inference from the truth of p , that one’s belief that p is true. But this result is incorrect, as the earlier perfectly ordinary example shows. So something has gone wrong in Zalabardo’s deployment of *PI*. Either the requirement is wrong in spirit, or something is wrong with the way in which Zalabardo is interpreting and handling its application in particular cases.

V

Let us take a closer look at *PI*. The requirement is hardly incidental to Zalabardo’s project. It plays a key role in his diagnosis of several problematic inferences that look *prima facie* acceptable on broadly foundationalist and reliabilist views of the sort he favors. For instance, why can’t the Moorean inference, “Here is a hand, so I am not a brain in a vat” yield knowledge that one isn’t a brain in a vat? Why can’t I learn from the newspaper that the Chicago Bulls won last night and then, by inference from the facts that the Bulls won and that the newspaper says they did, come to know that the newspaper report is veridical? Why can’t I learn on a number of occasions how much gas is in my car’s tank simply by reading the gauge and then, by inference from the record of what my gauge said on each occasion and how much gas was in the tank (as learned from the gauge), come to know that the gauge is reliable? In each case, Zalabardo maintains that a central failing is a violation of *PI*. In this way, Zalabardo hopes to use *PI* both to eliminate at one stroke a

number of problems that have plagued (among others) various foundationalist and externalist accounts of knowledge, thus safeguarding his own theory, and also to bolster his startling new skeptical argument.

Is *PI* correct? I am suspicious, but also unsure how to go about answering the question. Everything turns on what it would be for *S*'s believing *E* (evidence) to confirm *H* (hypothesis). And it is here that I begin to worry about how Zalabardo's theoretical apparatus is meant to connect with the world and with ordinary epistemic life. To bring out my worry, I will sketch what looks to me like a counterexample to *PI*. I will then consider whether a case could be made that it isn't a counterexample, but rather satisfies the requirement. I will in effect "play along" until I reach a point at which it becomes quite unclear to me what the theory really comes to.

Let's begin with what *PI* requires. Zalabardo glosses *PI* by saying, "My proposal is that what is wrong with these cases is that you wouldn't be less likely to believe your evidence if the hypothesis were false than if it were true" [p. 98]. If we take this gloss as our guide, it would appear that *PI* has straightforward counterexamples. Take, for instance, a version of Sosa's trash chute counter-example to Nozick's Sensitivity Requirement [Nozick (1999), pp. 145-6]. I open the trash chute on the twelfth floor of my apartment building, drop the trash bag in, and close the chute. Under ordinary conditions, I now know that my trash bag is somewhere in the bowels of the building. It seems that this knowledge would have to be inferential on Zalabardo's view: it is based, among other things, on the facts that I just dropped the trash into the chute, that the trash chute has worked impeccably for years, that nothing appears out of the ordinary and there are no warning signs or other indications of malfunction, etc. But of course it is compatible with all my evidence, though extremely improbable, that the bag snagged somewhere in the chute on the way down. And given that I deposited my bag in the chute in the ordinary way, it seems just as likely that I would believe all of my evidence if the trash bag were not now in the bowels in the building but had instead snagged on the way down, so I "wouldn't be less likely to believe [my] evidence if the hypothesis were false than if it were true". Straightforward counterexample, or so it seems – a counterexample that is just one instance of many such examples (from both ordinary life and science) in which we have inferential knowledge even though we would be just as likely to believe as we do even if the conclusion of our inference were false.²

Is this really a counterexample? If we look to Zalabardo's characterization of incremental confirmation, we are told that what *PI* requires is that the ratio of the following conditional probabilities be "sufficiently high":

$$P(S \text{ believes } E|H) / P(S \text{ believes } E|\sim H)$$

Perhaps then we should look more carefully at the relevant ratio of conditional probabilities. But how are we to determine the probability of my believing all my evidence conditional on the trash not now being in the bowels of the building? What are we to hold fixed in evaluating this? For instance, should we hold fixed that I did indeed put the trash into the chute in the ordinary way and closed the door? It seems plausible that we should. It's true that if I hadn't dropped my trash in the chute, it would be very improbable that my bag is now the basement and also very improbable that I would believe I had dropped my trash down the chute. But why should this fact have any bearing on *this* case? After all, what we want to know about is the relation between the location of my trash bag and my believing what I do, *when I have dropped my bag down the chute*. What I would believe when I don't even use the chute just seems irrelevant. But if in evaluating the relevant conditional probabilities we hold fixed that I put the trash in the chute and shut the door, then it seems that the probability of my believing E given H is pretty much the same as the probability of my believing E given not-H. So we still have a counterexample.

It might be argued that even if we hold fixed that the trash was dropped down the chute, certain other considerations would decrease the denominator of the ratio – so that this example would still satisfy *PI*. For instance, it might be claimed that if the trash is not now in the basement (because it snagged on the way down), then it is more likely for something noticeable to be out of the ordinary, so that I would be less likely to believe that nothing is out of the ordinary. It might likewise be claimed that if the trash is not now in the basement, then it is more likely that the chute would have malfunctioned in the past in a way that I would be aware of, so that I would be less likely to believe that the chute is reliable. But why can't we simply stipulate that in the world of the example, these possibilities are exceedingly unlikely? And why can't we simply stipulate that there will be similarly unlikely possibilities that will lower the value of the numerator correspondingly by lessening the conditional probability of my believing my evidence conditional on my trash being in the basement? For instance, there is the possibility that while dropping my trash into the chute, I have an amnesia attack and forget that I did so. Why can't we stipulate that in the world of the example, these and other such probabilities render the numerator and denominator roughly equal? It seems that I would know that my bag is now in the bowels of the building nonetheless. In fact, I have to admit that all of those very unlikely possibilities strike me as simply irrelevant to whether I know my trash bag is in the bowels of the building when I've dropped it down the chute in the ordinary sort of case.

Maybe some of these questions are to be dealt with by applying the Mixing Principle. For instance, via the Mixing Principle we could evaluate the probabilities of my believing my current evidence given that my trash bag is in the bowels of the building (i.e., $\text{prob}(\text{Believe } E|H)$) and of my believing

my current evidence given that it isn't (i.e., $\text{prob}(\text{Believe } E|\sim H)$) in a way that takes account of the probabilities of A1 that I dropped my bag down the chute, A2 that I accidentally left it in the elevator, A3 that I left it in my apartment, A4 that I didn't have a bag of trash at all and only hallucinated it, and so on and so on. To do this, we would need to assign probabilities to all of these possibilities. And now I have to lay my cards on the table. What on earth is supposed to determine the probabilities that we have been discussing? We are talking about a case in which I did drop my trash bag down the chute. With regard to that case, what are we to say about the objective probability that I didn't drop my trash bag down the chute? What is the objective probability that I left it in my apartment? What is the objective probability that there are warning signs not to use the chute, in the case in which there are in fact no such warning signs, the chute is in fact operating normally, no one thinks otherwise, etc.? Supposing such considerations to be relevant at all to whether I know where my trash bag is in this case (which I doubt), I have to admit that I am completely at sea with these questions. But if we are to perform the relevant calculations, we have to answer them.

What rules are to govern this game? *How* are we supposed to go about applying Zalabardo's particular notion of probability to this case, and *what* are we to go on in evaluating claims about categorical and conditional probabilities? We are told that what is at issue is an objective probability function crucially related to the laws of nature, that "the history of the universe" does not "play a role in the determination of probabilities," and that probabilities will be determined by "the laws [of nature] alone" [p.70], "with categorical probability ... defined as conditional probability on logical truth" [p. 73]. But these specifications vastly underdetermine how we might go about assigning probabilities, and it isn't always clear to me that Zalabardo's own declarations about probabilities even cohere with them. For instance, in the discussion of Kripke's counter-example to Nozick's theory, he writes, "The proposition that there are blue barn facades in the area (BLUEFAKES) is false, but its probability is increased by the presence of red barn facades in the area" [p. 122]. Really? What grounds such a claim? *Which* laws of nature have this effect, and how does Zalabardo know this? If the claim is just stipulative, then can we reframe the counterexample with different stipulations, so as to render it a genuine counterexample after all? Suppose, for instance, that the townspeople in fake-barn county (all good, law-abiding citizens) had passed a law that no barn facades are to be painted blue. Then would the presence of red barn facades in the area increase the probability that there are blue barn facades in the area? What if we stipulate that in the world of the example, the laws of nature are such that unlike red paints, blue-pigmented paints don't adhere to the materials from which the barn facades are constructed? I have to admit to some worry that Zalabardo's judgments about probabilities may

on occasion be driven by antecedent judgments about what confirms what, or about what the outcome should be in a given example.

So does my believing what I do in the usual case when I let go my bag and shut the chute confirm that my bag is now in the bowels of the building? It certainly seems that in a world such as ours, and given a psychology as generally reliable as mine is, my believing what I do would be good evidence that the bag is in the bowels of the building. If you know that I believe those things, and if you also know some other relevant stuff about the world, you can quite reasonably infer that my trash bag is now in the bowels of the building. So maybe *PI* is satisfied after all. But it isn't satisfied, if we understand it according to the instructions given in Zalabardo's gloss, since if the bag were not now in the bowels of the building, it would be stuck somewhere on the way down and I would be exactly as likely to believe what I do. Nor is *PI* satisfied on at least one way of understanding Zalabardo's specification of it in terms of his theory of incremental confirmation, namely, when we apply the theory holding fixed that I *did* drop the bag down the chute. And when we try to apply the theory taking account of other considerations, I at any rate quickly find myself at sea. So what are we to think? Perhaps the gloss on *PI*, and perhaps Zalabardo's account of incremental confirmation, don't track the notion of "good evidence" that I alluded to a moment ago? Or perhaps what we are given here, despite the appearance of precision, is not precise enough to be readily evaluable. I suspect that this is a danger one runs when one utilizes a highly developed mathematical theory without careful explanation of how it is to be interpreted and applied to particular cases.

VI

Once we notice that there are difficulties in the application of *PI*, the possibility opens up of thinking differently about the positive epistemic status of such beliefs as that we are not BIVs, not disembodied spirits being deceived by a Cartesian evil demon, and the like. After all, it would seem, *prima facie*, that we have a great deal of evidence against these hypotheses [Leite (2010), (2013)]. For instance, we currently lack the technology to create brains-in-vats, and there are no deceiving evil demons. Those considerations certainly seem to be decisive evidence against the BIV and Evil Demon hypotheses. So our beliefs in the non-obtaining of the skeptical hypotheses seem to be deeply evidentially enmeshed in our overall belief systems. We might say, perhaps exaggerating just a bit, "Everything tells against the skeptical hypotheses, and nothing is in their favor." Once *PI* is rejected (or at least, once it is ignored, given the unclarities about how it is to be applied), we can tell a story about the place of such beliefs in our overall belief system that both vindicates and explains their positive status. This would be a story

that, in a sense, locates them in the category of *inferential knowledge*, though we might think that for other reasons the simple Moorean argument and bootstrapping track-record arguments can't yield knowledge.

Such an approach should be unobjectionable, in broad brushstroke, to someone with the sorts of externalist and reliabilist sensibilities that Zalabardo shares. In particular, no worries about circularity or begging the question should arise, given Zalabardo's broad orientation [pp. 149-50]. Moreover, the fact that these evidential relationships are part of the story about how we know such things does not dictate that one could gain this knowledge for the first time through explicit, conscious inference from considerations about the external world; nothing I have said commits me to treating the Moorean inference as a good one or to saying (with the Moorean Dogmatist) that conscious, explicit reasoning from empirical facts about the external world could enable one to acceptably arrive for the first time at the belief that one is not a disembodied spirit being globally deceived by an evil demon. As I have argued elsewhere, requirements on epistemic responsibility would block one from coming to know such things in this way [Leite (2011b)]. So rather than looking for a single principle, such as *PI*, to do so much work at one fell stroke, it seems to me much more plausible – and much more what one would expect, given the variegation and multiplicity emphasized by Zalabardo's overall approach – that different considerations will have to be brought into play to explain what is going right and what is going wrong in various cases. Having taken this approach, we would be free both to explain – without appeal to “innate belief-forming predispositions” of dubious value – how we know that we are not BIVs or victims of a Cartesian evil demon, and also how we know in at least some cases that our beliefs are true.

Department of Philosophy
Indiana University, Bloomington
1033 E. Third St., Bloomington, IN 47405, USA
E-Mail: aleite@indiana.edu

NOTES

¹ This is just one sort of case with more or less this sort of structure; variation is possible particularly concerning the extent to which one is confident of one's view before double-checking.

² In a different context Zalabardo suggests [fn. 9, p. 129] that this case has the “same structure” as Vogel's ice cube counterexample to Nozick's tracking theory. A natural thought, then, would be to try to handle it analogously to Vogel's example. However, this treatment would be implausible, since unlike the melting of ice, there is no law of nature concerning what happens when one drops a trash bag down a trash chute.

REFERENCES

- LEITE, A. (2013), “‘That’s Not Evidence; it’s Not Even True!’”, *Philosophical Quarterly*, 63:250, pp. 81-14 (published online December 17, 2012)
- (2011a), “Austin on Dreams and Scepticism,” in M. Gustafsson and R. Sørli (eds.), *The Philosophy of J. L. Austin*, Oxford, Oxford University Press, pp. 78-113.
- (2011b), “Immediate Warrant, Epistemic Responsibility, and Moorean Dogmatism,” in *Reasons for Belief*, A. Reisner and A. Steglich-Petersen (eds.), Cambridge, Cambridge University Press, pp. 158-79.*
- (2010). “How to Take Skepticism Seriously,” *Philosophical Studies*, 148:1, pp. 39-60.
- SOSA, E. (1999), “How to Defeat Opposition to Moore,” *Philosophical Perspectives*, 13, *Epistemology*, edited by J. Tomberlin, Oxford: Blackwell.
- ZALABARDO, J. (2012), *Scepticism and Reliable Belief*, Oxford, Oxford University Press.

RESUMEN

Este artículo examina las propuestas de José Zalabardo sobre el conocimiento, el escepticismo sobre el mundo exterior, y nuestro conocimiento de que nuestras creencias son verdaderas. El principio PI de Zalabardo ocupa aquí un lugar central: “S puede tener conocimiento inferencial de *H* basado en la evidencia proporcionada por *E* si y sólo si la creencia de *S* en *E* confirma *H*”. Este principio parece ser vulnerable a algunos contraejemplos. Para decidir sobre el problema se requiere un examen de la teoría probabilista de la confirmación que ofrece Zalabardo y es aquí donde surgen las preguntas sobre cómo se intenta aplicar el aparato matemático de Zalabardo. Se ofrece una explicación alternativa de nuestro conocimiento de que no somos cerebros en cubetas así como una réplica al nuevo argumento escéptico de Zalabardo de que no podemos saber si nuestras creencias son verdaderas.

PALABRAS CLAVE: *escepticismo, confirmación, conocimiento inferencial, probabilidad, cerebros en cubetas.*

ABSTRACT

This paper scrutinizes José Zalabardo’s proposals regarding standing knowledge, external world skepticism, and our knowledge that our beliefs are true. Central here is Zalabardo’s principle PI, “*S* can have inferential knowledge of *H* based on the evidence provided by *E* only if *S*’s belief in *E* confirms *H*”. This principle appears vulnerable to counterexample. Adjudication of the issue requires scrutiny of Zalabardo’s probabilistic theory of confirmation, and here questions arise about how Zalabardo’s mathematical apparatus is meant to be applied. An alternative account of our knowledge that we are not brains in vats is offered, and a reply is offered to Zalabardo’s new skeptical argument that we cannot know that our beliefs are true.

KEYWORDS: *Skepticism, Confirmation, Inferential Knowledge, Probability, Brains in Vats.*