Article

CrossMark
click for updates

Bootstrapping Divine Foreknowledge? Comments on Fischer

Alan R. Rhoda

Christian Theological Seminary

John Martin Fischer's work has long been at the forefront of discussions on divine foreknowledge and human freedom. I have a great deal of respect for his work in this area and have learned much from him, so it's an honor to have been invited to comment on his recently published collection of essays on this topic (Fischer 2016). Given my space limitations and time constraints, I have decided to focus my comments on what he dubs the "Bootstrapping View" in section V of the volume's introductory essay. The view is a novel proposal for how God might be able to have infallible foreknowledge of the occurrences of causally indeterministic events. My comments on this proposal will be mostly expository, aimed at clarifying key terms and making important distinctions, but also critical. I do not think Fischer's proposal—or any proposal, for that matter—can reconcile infallible divine foreknowledge with future contingency. I believe this for fundamental reasons having to do with the nature of future contingency, knowledge, and truth.

In Part I will define future contingency and distinguish between different types of propositions about the future. This will help me to formulate Fischer's bootstrapping argument with greater precision. In Part II, I offer what I believe to be a faithful and charitable reconstruction of the bootstrapping argument and explain where I think it goes wrong. My main objection turns on a rather subtle point, namely, that concepts like *knowledge* and *truth* can be understood in either of two contrasting ways: (a) in a strict or

unqualified way, and (b) in a looser way that approximates to the strict conception but that isn't so demanding. I argue that the argument must ultimately equivocate because to succeed it must show that God can know certain future contingent propositions in the strict sense when, by the very nature of the case, those propositions cannot be strictly true, and therefore cannot be strictly known.

Future Contingency

In a footnote at the very beginning of section V, Fischer writes that he uses the term "future contingent"

> simply to mean a proposition that could (in some relevant sense, to be supplied by context) be or have been otherwise. I do not take it that "future contingent" implies indeterminism. I thus assume that there might be future contingents in a causally deterministic world (Fischer 2016, 31).

I find this characterization of future contingency to be misguided, for two main reasons. Discussion of these will allow me to clarify what I take future contingency to be and to introduce some terms and distinctions that will aid my analysis of Fischer's bootstrapping argument.

First, Fischer's characterization of a future contingent as a "proposition that could … be or have been otherwise" is too broad. As for how a proposition could be otherwise, I presume that he means otherwise *with*

Links
Researchers

respect to truth-values.[1] But if so, then it's not clear how *future* contingency on his view differs from *logical* contingency. And yet clearly the two should be distinguished, most obviously because logical contingency has no essential connection to the future, but also for dialectical reasons. After all, the problem of foreknowledge and future contingency has a very long history, and the notion of future contingency that lies at the heart of the problem is much narrower than Fischer's characterization. We can see this as follows. First, while virtually no one has thought it problematic that God could have infallible knowledge of logically contingent truths about the past and present, many have thought there is problem with God's having infallible knowledge of contingent truths *about the future*. Second, and more specifically, with respect to the future, no one doubts that God could have infallible knowledge of causally determined future events. Hence the problem centers on God's knowledge of a *causally indeterministic* future. That's the sort of future contingency that generates the problem, and so, despite what Fischer says in this note, the broader dialectical context suggests that future contingents *should* be understood in a manner that entails causal indeterminism. This is especially so in the narrower context of section V of Fischer's introduction where the explicit purpose is to defend the possibility of divine foreknowledge of future contingents in a causally indeterministic world. Accordingly, in what follows, I will take future contingency to entail causal indeterminism. This won't have any bearing on my evaluation of Fischer's bootstrapping argument, but it is important to be clear on how key terms are being used.

Second, it is better to think of future contingents as causally possible *events* rather than as propositions. This is partly because it is events, especially human free choices, that we wish to reconcile, if possible, with divine foreknowledge. God's foreknowledge presumably consists (in part) in his knowing propositions *about* future contingents. I'll call these *future contingent propositions* (FCPs). But the contingencies that matter for us lie in the concrete world of events, not in the abstract realm of propositions.

Another reason for not taking future contingents to be propositions is that membership in the class of FCPs is unstable. As Yoda says in *The Empire Strikes Back*, "Always in motion is the future." <Caesar will cross the Rubicon in 49 B.C.>[2] may have been an FCP in 50 B.C., but it isn't one now. The mere passage of time ensures that no proposition about finitely distant events is always about the future, much less the contingent future. Future contingencies may also be foreclosed by prior contingencies. Suppose that before persons A and B say "I do" at the marriage altar, it is a future contingent whether they get married. If, however, C should murder A prior to any such marriage, then that contingency will have been foreclosed by A's death, with the result that <A and B will marry> is no longer an FCP (barring the possibility of A's miraculous resurrection).

What we have on the one hand are causally possible events, and on the other hand propositions that purport to describe those events. Some of these propositions may be *about* future contingents (FCPs), but only if there *are* future contingents, and that's not something we can discover by inspecting the propositions. We may *suppose* <Jones will mow his lawn on Wednesday> to be an FCP, but the truth or falsity of that supposition depends on the actual current causal matrix of reality. If thoroughgoing determinists are right, then no propositions succeed in being about future contingents except in the trivial sense that <There are no unicorns> succeeds in being about unicorns. That is, while there would still be propositions that employ the *concept* of a future contingent, the extension of the concept would be empty.

Future contingent propositions or FCPs are a subclass of propositions about the future, which fall into two main groups. Members of the first group represent the future occurrence of an event as a settled fact about some putative actual future. These propositions are commonly expressed in terms of what "will", "will not", or "is going to" happen. I call them *settled future propositions* (SFPs). Members of the second group represent the future occurrence of an event as something yet-to-be-settled, thus leaving it an *open question* whether the event occurs in the future or not. These propositions are commonly expressed in terms of what "may", "might", "could", or "will probably" happen. I call them *open future propositions* (OFPs). Concerning a coin about to be tossed, for example, <The coin will land heads> is a SFP and <The coin might and might not land heads> is an OFP. When terms like "might" and "probably" in an OFP are understood non-epistemically to connote the causal possibility or probability of an event's occurring, and are conjoined in a manner that implies the event is both causally possible and not causally necessary (e.g., "might and

might not"), then these propositions become especially apt vehicles for describing future contingents.

Overlaying the class of future contingent propositions with the SFP/OFP distinction yields a further distinction that will be useful when analyzing the bootstrapping argument. Let's call FCPs that are SFPs *settled future contingent propositions* (SFCPs) and FCPs that are OFPs *open future contingent propositions* (OFCPs). For example, if a coin flipping scenario is genuinely indeterministic, then <The coin will land heads> is both an SFP and an SFCP. Without the indeterminism it's just an SFP.

Finally, to evaluate whether a proposition about the future is a future contingent proposition (FCP) and therefore either an SFCP or OFCP rather than merely an SFP or OFP, we need a definition of future contingency. I offer two such definitions. Both reflect the stipulation, argued for above, that future contingency entails causal indeterminism. Thus, if we take a "future" to be a complete, linear extension of the actual past and present, then we can define a future contingent as an event that occurs in at least one causally possible future, but not in all causally possible futures. An alternative, extensionally equivalent definition can be given using the notion of single-case objective probabilities or *chances*. Since a causally determined event E has every chance of occurring, i.e., $Ch(E) = 1$, and a causally impossible event has no chance of occurring, i.e., $Ch(E) = 0$, a future contingent must be an event whose chance of occurring in the future is both greater than zero and less than one.

The Bootstrapping Argument

I turn now to an examination of Fischer's bootstrapping argument. Its stated aim is to establish the possibility of God's knowing *with certainty* on Monday a proposition like <Jones will mow his lawn on Wednesday> given an indeterministic world Fischer 2016, 37). Since the proposition in question is assumed to be an SFCP, we can restate the aim of the argument more generally as establishing the possibility of God's knowing *an SFCP* with certainty. The term "certainty", though, is ambiguous. It can mean *subjective* certainty (e.g., "I am certain that Jones will mow his lawn on Wednesday") where it describes the *credence*, degree of confidence, or strength of belief that a subject has in a proposition. Or it can mean *objective* certainty (e.g., "It is certain that Jones mows his

lawn on Wednesday"), where it describes the absence of any *chance* that things turn out otherwise.[3] So does Fischer mean to argue that God can know an SFCP with subjective certainty (i.e., believe it with maximal credence), with objective certainty (i.e., know it infallibly, without any chance of error), or both? Given the contrast he wants to draw between divine knowledge and *fallible* human knowledge, and since one can have subjective certainty in something that turns out to be false, I think that he must mean that God can know an SFCP with *at least* objective certainty, if not with both objective and subjective certainty.

The gist of Fischer's argument is reasonably clear in the text: (a) Use reflections on human knowledge to argue for the possibility of God's having at least *fallible* knowledge of an SFCP, and then (b) leverage divine omniscience to bootstrap that into *infallible* knowledge. Fischer, unfortunately, doesn't make the full structure of his argument explicit by detailing premises and intermediate conclusions. It took me a bit of work to piece it all together in a way that seemed both logically valid and argumentatively cogent. I'm not 100% sure that I have it exactly right, but with the possible exception of (6), I believe there is clear textual support for all of the other premises.

1. In general, the conditions necessary for being in a knowledge conferring situation (KCS) are consistent with fallibilism. (Assumption)
2. For any epistemic agent S, S knows that *p* if and only if S's epistemic situation with respect to *p* satisfies all the conditions necessary for being in a KCS, S believes that *p*, and *p* is true. (Definition of KCS; Unity of knowledge)
3. Therefore, if *p* is true then any epistemic agent who believes *p* and is in a KCS with respect to *p* has at least fallible knowledge that *p*. (From 1, 2)
4. Humans sometimes believe and are in a KCS with respect to a true SFCP. (Assumption)
5. Therefore, humans sometimes have at least fallible knowledge of an SFCP. (From 3, 4)
6. Anything about the future that humans know, God knows at least as well. (Assumption)
7. Therefore, God sometimes has at least fallible knowledge of an SFCP. (From 5, 6)
8. God knows *with objective certainty* whatever He believes. (Essential omniscience)
9. Therefore, God sometimes knows an SFCP *with objective certainty*. (From 2, 7, 8)

The expression "knowledge conferring situation" (KCS) is Fischer's noncommittal catchphrase for whatever must be added to belief and truth to get knowledge. (1) says that the conditions for being in a KCS (with respect to a given proposition) are consistent with fallibilism. For the argument to work, this need not be true for all propositions. Perhaps there are some that can only be known infallibly, or not at all. But it does have to be true with respect to SF-CPs that are epistemically accessible to humans, and this seems highly plausible. Fischer briefly canvasses a diverse sampling of theories of knowledge, all of which are compatible with fallibilism (Fischer 2016, 34–36). A *fallibilist* theory of knowledge is one that allows S sometimes to know that *p* in cases where S's epistemic situation *fails to guarantee* that *p* is true. In other words, in cases of fallible knowledge, S believes that *p* and is in a KCS with respect to *p*, but whether S knows that *p* depends on the further *external* condition that *p* is true. In contrast, in cases of *infallible* knowledge, S's epistemic situation with respect to *p* is truth-entailing, such that S's believing *p* and being in a KCS with respect to *p* is enough to guarantee that *p* is true.

Premise (2) spells out necessary and sufficient conditions for knowledge: being in a KCS with respect to given proposition, believing that proposition, and its being true. The universal quantification over epistemic agents is intended to capture Fischer's view that knowledge is "unified", by which he means roughly that the basic rules of knowledge work the same way for all epistemic agents, especially (in this context) humans and God (Fischer 2016, 35–36).

Given (1) and (2) it follows that any epistemic agent who meets all the conditions for knowledge that *p* has *at least* fallible knowledge that *p*. To be in a KCS that is consistent with fallibilism doesn't rule out the possibility of having infallible knowledge, but it doesn't give us license to infer that any knowledge that results from that KCS is infallible.

(1), (2), and (3) all seem right to me, but I should mention now, because of its relevance for my critical evaluation of the argument, that the concept of *knowledge* is arguably less straightforward than Fischer seems to think. He doesn't explicitly say so, but I think he believes the concept of knowledge to be *univocal*, i.e., such that it carries the very same sense across all instances of knowledge. Perhaps that's not

his view, but in any case, I submit that knowledge is an *analogical* concept, one that comes in degrees, and that subdivides into qualitatively distinct tiers, among which we should distinguish unreflective (animal) knowledge, reflective (human) knowledge, and perfect (divine) knowledge. One reason for thinking this is that nearly all the key concepts that figure prominently in analyses of knowledge plausibly admit of degrees, including truth (verisimilitude), belief (credence), reliability, justification, and warrant. And so, it's at least initially plausible that knowledge should admit of degrees too. Another reason for thinking knowledge is analogical is the familiar fact that one person can seemingly *know* the same proposition much more fully than someone else. The student who has memorized the Central Limit Theorem from a math textbook and can apply it correctly to basic textbook problems may be said to "know" the theorem, and on most theories of knowledge he does. But the math professor who has constructed five independent proofs of the theorem and can apply it in novel and sophisticated ways really *knows* it. Relative to the student, we might even say that he "superknows" it. The quality of his knowledge contains all the positive qualities of the student's knowledge, and much more. Likewise, relative to the professor, God superknows the theorem even further, to the superlative degree, for He instantly and with perfect clarity grasps *every* possible proof of the theorem, *every* logical consequence of it, and *every* possible application of it.

Much more needs to be said to flesh out and defend an analogical conception of knowledge—a project I hope to take on some day. But if I'm right, then the relevance for (1) and (2) is that it requires us to consider *what degree* or grade of knowledge a given KCS can confer. Presumably some lower-quality KCSs can only confer only low (e.g., unreflective) or medium (e.g., reflective) grades of knowledge. But perfect or divine knowledge presumably requires an exceptionally high-quality sort of KCS. Could this mean that God might know *less* of some matters than humans do because the requirements for divine-quality knowledge are higher? No, the issue here is quality, not quantity. Anything humans can know, God can know at least as well.[4] If humans can know *p* fallibly, then God can also know *p* fallibly. But such mundane-level knowledge doesn't entail divine-quality superknowledge, which entails knowing *p* infallibly. God's being able to know *p* fallibly, in a human-quality way, *doesn't* mean He can also know *p* infallibly, in a divine-quality way.

On the other hand, it doesn't mean He *can't* also know *p* infallibly in virtue of being in a higher-quality KCS. Turning now to (4), it says that humans sometimes satisfy all the conditions for *knowing* an SFCP. But do they? Setting the truth condition aside for now, is it true that humans sometimes believe and are in a KCS with respect to a SFCP? Fischer seems to answer in the affirmative. Based on information about Jones's beliefs, desires, and habits, he thinks we can have fallible knowledge on Monday of <Jones will mow his lawn on Wednesday> (Fischer 2016, 32). But while that proposition is clearly an SFP, it's not clear that it's also an SFCP, a proposition genuinely about a future contingent event. Whether the event described is a future contingent or not is a matter that depends on the causal landscape of reality. So let's *stipulate* that the causal landscape is indeterministic in this respect. Let's stipulate that, in relation to Monday, there are causally possible futures in which Jones mows on Wednesday and causally possible futures in which Jones does not mow on Wednesday. But now it must be asked: What is the ratio of mowing futures to non-mowing futures? And how does that bear on the possibility of reasonable belief in and knowledge of <Jones will mow his lawn on Wednesday>?

Enter Lewis's well-known *Principal Principle*, which articulates a commonsense rule for assigning rational degrees of confidence or credences to chance events. Simply stated, the Principle says that one's credence should match what one believes about the relevant chances (Lewis 1986). For example, if I believe that a normal six-sided die is perfectly fair, then I believe that the chance of a toss landing on any given face is 1/6. Hence, by the Principal Principle, my credence that a toss will land on 3 should also be 1/6.

Now a credence of 1/6 is obviously far too low for *knowledge* that the die's toss will land on 3. Indeed, it is far too low even for *belief* that it will land on 3. Based on my assumptions about the chances I should have a credence of 5/6 that it will *not* land on 3. Hence, *that's* the rational thing to believe. Minimally, belief requires a credence somewhat higher than 0.5. If you think the chance of some event is roughly equal to or less than 0.5, then you shouldn't believe that it will happen. If anything, you should believe the opposite, or withhold belief altogether.[5] As for knowledge, that seems to require a credence *much higher* than 0.5, and perhaps very close to 1. After all, even with a credence of 5/6 that the die won't land on 3, I don't think it'd be

right for me to claim to *know* that it won't land on 3.

Combining the preceding observations about belief and knowledge with the Principal Principle, it follows that *knowledge* that <Jones will mow his lawn on Wednesday> is only possible if our assessment of the chance of him doing so is very high. That means we would need to have evidence that the circumstances and his behavioral dispositions are tending *very strongly* in that direction. Generalizing, it seems that only a small subset of SFCPs are plausible candidates for even fallible knowledge.

But knowledge requires still more than high credence. It also requires that the proposition known be *true*. When it comes to SFCPs, this is problematic. The proposition <Jones will mow his lawn on Wednesday> is an SFP. As such, it represents it as a *settled fact* on Monday that Jones mows on Wednesday. Given that this event is indeterministic, however, it follows that it is *not* a settled fact on Monday that Jones mows on Wednesday because it remains causally possible that he does *not* mow. Obviously, it can't at the same time (Monday) be both an open-question and a settled fact how things turn out with respect to Jones's mowing on Wednesday. Since a proposition is true if and only if it corresponds to reality, and since <Jones will mow his lawn on Wednesday> misrepresents reality by presenting something as settled when it isn't, the proposition isn't true, and hence not even a possible candidate for knowledge, fallible or otherwise.[6]

If that's the full story on truth, then the bootstrapping argument is dead in the water, even before we get to the bootstrapping part. There is, however, a way to salvage the situation and defend (4). It lies in recognizing, as I suggested above, that knowledge and truth are analogical concepts. Both admit of degrees,[7] and both have an intrinsic maximum consisting of perfect or divine-quality knowledge, in the one case, and perfect correspondence with reality, in the other. Our ordinary usage of words like "know" and "true" is thus like our ordinary usage of descriptive terms like "flat".[8] We commonly *say* of things like tables that they are "flat" and of propositions like <This table is flat> that they are "true" and of persons who make such claims that they "know" what they're talking about. And we *say* things like this all the while knowing that a microscopic examination of the table's surface would clearly reveal that it is *not flat*, that a proposition describing it as flat is *not true*, and that knowledge claims to

the contrary are, at the very least, *inexact*. To make sense of this linguistic behavior, we need to recognize a distinction between *loose* and *strict* ways of speaking. We can "know" (loosely speaking) and it can be "true" (loosely speaking) that the table is "flat" (loosely speaking) because the table is *flat enough* for the practical contexts in which such claims are typically made. Likewise, even though SFCPs like <Jones will mow his lawn on Wednesday> are all *strictly false* and thus not candidates for *knowledge*, strictly speaking, they can nevertheless be *true enough* if the chance that events do not turn out as expected is small enough to be practically negligible. And they can be *knowable enough* if one's KCS secures a tight enough connection with the truth (strictly speaking) that only a small and practically negligible chance for error remains.

In short, my proposal is that while (4) and by extension (5) are strictly false—SFCPs are neither strictly true nor strictly knowable—they are nevertheless acceptable if (and only if) "true" in (4) and "knowledge" in (5) are understood in accordance with the somewhat loose usage of ordinary language.

Moving on, (6) and (7) are easy to defend. Regarding (6), God obviously knows the future at least as well as any human does because God has immediate access to *all* information about the past and present that in any way bears upon the future, far more information than all humans put together possess. And (7) follows from (6) and (5). Given what I've just said in defense of (4) and (5), however, the term "knowledge" in (7) must be understood in a somewhat loose sense. It can't be knowledge strictly and unqualifiedly understood because, after all, the proposition in question isn't strictly true.

Now here's where the bootstrapping operation kicks in. Fischer proposes that God's essential omniscience can elevate fallible knowledge into infallible knowledge. According to (8), because God is essentially omniscient, God knows with *objective certainty* whatever He believes. This means that if God believes *p* then there is zero chance of *p*'s turning out false. Essential omniscience thus entails divine infallibility. But then (9) seems to follow from (2), (7), and (8). By (7) God sometimes has *at least* fallible knowledge of an SFCP. By (2), if He has knowledge (fallible or otherwise) of an SFCP, then the proposition is true, God believes it, and He is in a KCS with respect to it. By (8) if God believes it, then He knows it with objective certainty

(i.e., infallibly). Hence, God sometimes has infallible knowledge of an SFCP. God's believing the proposition and being in a KCS with respect to it thus suffices to *guarantee* that the proposition is true, which satisfies the requirements for infallible knowledge.

However, even though (9) may seem to follow, there is a problem with (8). As I argued above, *all* SFCPs are *strictly false* because they misrepresent an event as a *settled fact* when, as a future contingent, its occurrence remains an *open question*. I argued with respect to (4) that SFCPs could still be "true" and objects of fallible "knowledge", so to speak, provided those quoted terms are understood in accordance with the somewhat lax usage of ordinary language. But divine knowledge is absolute and perfect. Essential omniscience doesn't just mean that, necessarily, whatever God believes is "true", it means that, necessarily, God knows and therefore believes things to be *exactly* as they are. Hence, there can be no semantic looseness when it comes to divine-quality knowledge. What God divinely knows to be true must be *strictly true*, not merely "true enough". Likewise, what God *believes* must be strictly true as well. Or at least, it must be strictly true unless "believe" like "know" and "true" is used in a somewhat loose manner that admits credences less than one. But if "believe" is understood loosely, then (8) is false, for it is not the case that God knows with objective certainty whatever He "believes". Rather, God knows with objective certainty whatever He believes *with subjective certainty* (i.e., with maximal credence). And by the Principle Principle God can't believe with maximal credence something that He knows has less than maximal chance of obtaining. So God doesn't and can't believe with maximal credence an SFCP like <Jones will mow his lawn on Wednesday>. What he can believe with maximal credence and infallibly know is something more like <Jones will *probably* mow his lawn on Wednesday>. But that's an OFCP, not an SFCP. Hence, the argument fails to establish that God can know an SFCP with objective certainty.

To sum up, the heart of my critique of Fischer's bootstrapping argument can be stated as a dilemma. Either we are prepared follow ordinary language using terms like "knowledge", "truth", and "belief" somewhat loosely, or we aren't. If we are, then (1)–(7) can all be consistently understood in a way that is plausibly true. That is, in ordinary language contexts it may be a pragmatically close enough approximation to the

strict and unqualified truth to say that humans and God can "believe" and "know" that an SFCP is "true". But once we start invoking epistemic perfection concepts like essential omniscience, objective certainty, and infallibility, we shift to a theoretical context where such "loose use" is no longer admissible. In short, on this horn of the dilemma, the bootstrapping phase of the argument equivocates because it slides between loose and strict understandings of the key epistemic terms. As for the other horn, if we try to circumvent equivocation by stipulating that all key terms be understood in a strict and unqualified manner throughout, then premises (1) and (4) become false because SFCPs are strictly false and thus not possible objects of knowledge, strictly speaking. It follows that (5) and (7) are strictly false as well and that neither humans nor God can *know* an SFCP.

A Closing Reflection

My discussion has centered on Fischer's bootstrapping argument, and I have contended that it fails to show that God can have infallible knowledge of SFCPs. Might some other argument be able to accomplish what bootstrapping has not? No. Not if I am correct in saying that SFCPs are strictly false precisely because they represent the future as *settled* in some respect when, by virtue of the fact that we're dealing with future contingents, it *isn't* settled in that respect. In short, the very notion of a strictly true SFCP is incoherent. And since what is not strictly true cannot be strictly known, the very idea of infallible divine foreknowledge of an SFCP is also incoherent. The considerations I have raised in this essay therefore generalize well beyond the bootstrapping argument.[9]

References

- Fischer, John Martin. 2016. Our fate: Essays on God and free will. Oxford: Oxford Univ. Press.
- Lewis, David. 1986. "A subjectivist's guide to objective chance", in Philosophical papers, volume II, pp. 83–113. Oxford: Oxford Univ. Press.
- Smith, Nicholas J. J. 2008. Vagueness and degrees of truth. Oxford: Oxford Univ. Press.
- Unger, Peter. 2002. Ignorance: A case for skepticism. Oxford: Oxford Univ. Press.
-

End Notes

[1] Obviously, a proposition can't be otherwise with respect to its *meaning*, for a proposition just *is* its meaning.

[2] I use angle brackets <> enclosing a declarative sentence to denote *propositions*. Double-quotes "" around a clause denotes the sentence quoted, not a proposition.

[3] Note the shift from "will mow" to "mows". Whereas credences apply primarily to propositions, chances apply primarily to events. The event in question is not *Jones's going to be mowing on Wednesday*, but *Jones's mowing on Wednesday*.

[4] With the possible exception of first-person *de se* knowledge. Arguably, only John Fischer can know <I am John Fischer>.

[5] You might *wager* that a low-chance event will happen, and choose to accept the risk that comes with it, but betting is not the same as believing.

[6] The problem remains even if we suppose an eternalist ontology of time, according to which a unique and complete series of future events exists. Such an ontology would render it true *simpliciter* that <Jones mows his lawn on Wednesday>, but it wouldn't render it true *on Monday* that <Jones will mow his lawn on Wednesday>. Truth *simpliciter* is evaluated from the so-called "God's eye" perspective which, on eternalism, includes (we may suppose) the event of Jones's mowing. But truth-*at-a-time* is evaluated from the perspective of the time in question. And if the event in question is a future contingent with respect to that time, then there is no settled fact *as of that time* that the event is going to occur.

[7] For a detailed defense of the idea that truth comes in degrees, see Smith (2008).

[8] Peter Unger famously compares the predicates "knows" and "is flat" in the interests of defending a very thoroughgoing form of skepticism (Unger 2002, esp. ch. 2). I am interested here in pressing the comparison in the opposite direction, toward the possibility of fallible knowledge, by recognizing that "knows", like "flat", is subject to loose use.

[9] I am grateful to my colleague David Stout for helpful feedback on a draft of this paper.