# Reliabilism and Imprecise Credences

## 1 Introduction

Are there situations in which imprecise credences are rationally required? The following example from Sturgeon (2010) is meant to show that the answer is 'yes':

> **(Black Box)** A black box sits in front of us, and we're rationally sure that it contains many balls that have been thoroughly mixed; further, we're 'rationally sure solely that exactly 80% to 90% of balls in the box are red' (*ibid.*, 129-130). Given that we've no other information about the contents of the box, how confident should we be that a ball drawn randomly from it will be red?

Sturgeon thinks that we 'should be exactly 80% to 90% sure' (*ibid.*). Given that the relevant evidence is rough or non-specific, he thinks that our credence in RED—the proposition that we'll draw a red ball at random—shouldn't be precise. Instead, our credence should be imprecise, where a subject's imprecise credences are represented by a set of credence functions, or what van Fraassen (1990) calls a *representor* (347). In particular, Sturgeon thinks that our imprecise credence should be spread over the interval [0.8, 0.9]—for any $x$ in the interval, our representor should contain a credence function that assigns $x$ to RED.

The status of imprecise credences is up for debate. For instance, like Sturgeon, Hájek and Smithson (2012), Joyce (2010), Levi (1974), and Walley (1991) hold that some situations require imprecise credences—or *imprecision* for short. But White (2010) and Elga (2010) deny this. In addition, Carr (ms) argues that the reasons put forward for having an imprecise credence in $p$ might in fact be reasons for higher-order uncertainty as to what precise credence to assign $p$. And Schoenfield (2017) suggests that certain situations in which imprecision seems to be required are really ones in which it's indeterminate what precise credence one has or ought to have (680). Relatedly, Levinstein (2019) countenances imprecision, but thinks that there's an interpretation of imprecision on which for one to have an imprecise credence is just for it to be indeterminate what precise credence one has.[1]

In this paper, I'll grant Sturgeon's intuition about Black Box. And I'm not concerned primarily with whether friends or foes of imprecise credences are right. Instead, I'll address a relatively neglected question. Suppose that some situations require imprecision. I take it, then, that sometimes imprecise credences are justified whereas precise credences are unjustified. But *what is it for imprecise credences to be justified in the first place*?

I won't attempt to answer the question in full. Instead, I'll focus on how a *reliabilist* might answer the question. I'll also assume, at least initially, that to have an imprecise credence in $p$ is to have a determinate first-order doxastic attitude towards $p$. I'll assume for the sake of argument that some situations require imprecision—and that these are not merely situations in which one ought to be uncertain about what precise credence to have or situations in which it's indeterminate what precise credence one has or ought to have.[2] But towards the end of the paper, I'll reconsider the question with respect to indeterminate credences—or to Levinstein's interpretation of imprecise credences.

Friends of imprecise credences—or *imprecisers* for short—might think that there's

[1]A word on terminology: Levi (1985) reserves the term 'imprecision' for cases in which we have a precise credence but are unable to determine its exact value. This is *not* the notion of imprecision I have in mind when I talk about imprecise credences. I use the term 'imprecise' to talk about the lack of sharpness in an agent's credal state itself. To talk about such lack of sharpness, Levi, as well as Hájek and Smithson (2012), use the term 'indeterminate'. But I'll use the term 'indeterminate credence' to refer specifically to what Levinstein means by 'imprecise credence'.

[2]For example, on this view, when we say that one is required to have a credence of [0.8, 0.9] in $p$, the interval [0.8, 0.9] is not meant to represent uncertainty about or indeterminacy in one's credal attitude. Instead, it's supposed to determinately single out a particular, interval-valued credal attitiude.

no urgent need to answer the question above. They might think that its answer just depends on how a well-trodden issue in epistemology plays out—namely, that of which theory of doxastic justification, be it reliabilism, or evidentialism, or some other theory, is correct. In particular, they might think that we can be neutral about this issue while holding that there are justified imprecise credences; after all, one may hold that there are justified binary beliefs without having a settled view on what it is for a binary belief to be justified.

But matters are not that straightforward. As we'll see, various attempts at giving a reliabilist treatment of imprecision will prove problematic. And if it turns out that reliabilism is incompatible with imprecision, that would be significant. For it would turn out that imprecisers cannot be neutral as to which theory of justification is correct: an impreciser would have to reject reliabilism. Further, a reliabilist would have to reject imprecision.

Here's the plan. In section 2, I'll present four reliabilist accounts of justified precise credence. But my aim is not to assess these accounts on their own merits.[3] Instead, I want to explore whether we may extend them to deal with imprecision. To anticipate, I'll consider various such extensions in section 3 but argue that they face problems. I'll conclude by exploring how reliabilists may accommodate indeterminate credences—even if they cannot accommodate imprecise credences understood as determinate first-order attitudes.

## 2 Four Reliability Theories of Justified (Precise) Credence

One might think that reliabilism isn't well suited to accounting for imprecise credences because it isn't well suited to accounting for credences in general, precise or otherwise. After all, reliabilists typically spell out doxastic justification in terms of truth-conduciveness. But credences, unlike all-or-nothing beliefs, do not in general

[3]For more in-depth discussion of reliabilist theories of justified (precise) credence, see Dunn (2015), Tang (2016), and Pettigrew (forthcoming). Also, since my focus is on whether reliabilism is compatible with imprecision, and not on giving a full-fledged defence of it, I'll set aside the usual problems that reliabilism is thought to face, such as the generality problem and the New Evil Demon Problem.

admit of truth or falsity. For instance, a credence of 0.5 that it's raining is not the kind of state that is either true nor false. Yet, like beliefs, credences may be justified or unjustified. For example, our credence of 0.5 that it's raining is unjustified if we've very strong evidence for rain.

But the above is not my reason for worrying about whether reliabilism is compatible with imprecision. It would be too hasty to claim that, because credences are not in general truth-apt, reliabilists cannot account for their justifiedness. In fact, there are a number of ways in which reliabilists—be they *process* reliabilists or *indicator* reliabilists—might spell out the justifiedness of precise credences.

### 2.1 An appeal to objective probability

According to traditional process reliabilism, a belief is justified if and only if it's produced by a reliable process—that is, one that tends to produce a high ratio of true to false beliefs. Granted, credences are not in general true or false. But there might be an analogue of truth for credence. In particular, one might think that while beliefs aim at truth, credences aim at objective probabilities—that just as there's something epistemically good about beliefs being true, there's something epistemically good about a credence in $p$ matching the objective probability of $p$'s being true.[4] With this thought in mind, consider:

**(Probability)** For any $x \in [0, 1]$ and any $p$, a credence of $x$ in $p$ is justified if and only if it's produced by a process that tends to produce a high proportion of credences that match (or at least approximate) the corresponding objective probabilities.[5]

Whereas traditional process reliabilism spells out justified belief in terms of a high proportion of beliefs matching the truth, (Probability) spells out justified credence in terms of a high proportion of credences matching the corresponding

[4] Hájek (ms) argues that whereas belief is vindicated by truth, credence is vindicated by objective chance. Carr (2016) also offers a few reasons for such a view.

[5] I discuss this theory in Tang (2016) and argue that it faces problems qua theory of justified precise credence. But here, I'll ignore such problems—my focus is on whether there's a natural way to extend the theory to deal with imprecision.

objective probabilities. Note that invoking different notions of objective probability will give us different versions of (Probability). Also, instead of requiring that credences match the relevant objective probabilities exactly, we may require only that the match be approximate. Given that our credences will hardly match the relevant objective probabilities exactly, the former requirement would make it too hard for them to be justified.

### 2.2 An appeal to calibration

A second option appeals to the notion of *calibration*, illustrated by van Fraassen (1984) as follows:

> [C]onsider a weather forecaster who says in the morning that the probability of rain equals 0.8. That day it either rains or does not. How good a forecaster is he? Clearly to evaluate him we must look at his performance over a longer period of time. Calibration is a measure of agreement between judgments and actual frequencies [...]. This forecaster was perfectly calibrated over the past year, for example, if, for every number $r$, the proportion of rainy days among those days on which he announced probability $r$ for rain, equalled $r$. (245)

According to van Fraassen, a good weather forecaster is calibrated. If you think that a good belief-forming process should likewise be calibrated, you might like the following theory of justified credence:

**(Calibration)** For any $x \in [0, 1]$ and any $p$, a credence of $x$ in $p$ is justified if and only if it's produced by a calibrated (or approximately calibrated) process, where a process is calibrated (or approximately calibrated) if and only if, for any $y \in [0, 1]$, $100y\%$ (or approximately $100y\%$) of the propositions in which the process causes us to have a credence of $y$ are true.[6]

[6] Versions of (Calibration) have been put forward by Lam (2011), Goldman (2012), and Goldman (1986) (212-215; 26; 113-115). See also Dunn (2015) and Tang (2016).

Just as traditional process reliabilism doesn't demand that a process be perfectly reliable for it to produce justified beliefs, (Calibration) doesn't demand that a process be perfectly calibrated for it to produce justified credences. Also, instead of requiring that a process be calibrated, we may require that it be *potentially* calibrated (van Fraassen 1983, 302-305). Whereas actual calibration measures agreement between credences and actual frequencies, potential calibration measures agreement between credences and hypothetical frequencies. For example, suppose a belief-forming process comes into existence, produces a credence of 0.7 in $p$, and then stops producing any more credences. This process is not calibrated insofar as we're concerned with actual calibration. But it might still be potentially calibrated. For instance, if it had gone on to produce more credences of 0.7 in other propositions, it might be that 70% of the relevant propositions would have been true.

### 2.3 An appeal to accuracy

Instead of appealing to calibration or to matches between credences and objective probabilities, we might appeal to the notion of *accuracy*. The intuitive idea is that the higher our credence in a truth or the lower our credence in a falsehood, the more accurate it is. And just as there seems to be something epistemically good about having true beliefs and avoiding false ones, there seems to be something epistemically good about having accurate credences and avoiding inaccurate ones.

Following Joyce (1998), we may measure the (in)accuracy of a credence of $x$ in $p$ with the Brier score, given by the formula $(x - T(p))^2$, where $T(p)$ equals 1 if $p$ is true and 0 if $p$ is false.[7] The lower the Brier score of a credence, the more accurate it is, and the higher its Brier score, the more inaccurate it is. For example, our credence of $x$ in $p$ is perfectly accurate when $x = 0$ and $p$ is false (in which case its Brier score is 0). It's perfectly inaccurate when $x = 1$ and $p$ is false (in which case its Brier score is 1). Appealing to Brier scoring, here's a corresponding theory of

[7]The Brier score is due to Glenn W. Brier (1950), who first proposed that we use it to gauge the accuracy of weather forecasts.

justified credence:

**(Brier)** For any $x \in [0, 1]$ and any $p$, a credence of $x$ in $p$ is justified if and only if it's produced by a process that tends to produce credences with a low average Brier score.[8]

Instead of appealing to Brier scoring, we might appeal to other scoring rules. Doing so should lead to theories that, though different from (Brier), are still similar in spirit to it.[9]

### 2.4 An appeal to reliable indication

Process reliabilists might find the three theories above congenial. But an *indicator* reliabilist might also provide us with a theory of justified credence. According to William Alston, a belief is justified if and only if it's based on a ground that's sufficiently *indicative* of the belief's truth. But what is it for a belief to be based on a ground, and what is it for a ground to be indicative of the belief's truth?

According to Alston, a belief $b$ is based on a certain ground $g$ just in case $g$ is the input to the belief-forming mechanism responsible for producing $b$ as its output, where such an input takes the form of 'something psychological—some psychological state or process' such as a belief, a memory, or an experience (Alston 2005, 126-127; 83).[10] For example, if we believe that there's a cup on the table, and this belief is the output of a belief-forming mechanism that takes our visual experience as of there being a cup on the table as its input, then such a visual experience is the ground on which our belief is based.

Further, Alston takes a ground to sufficiently indicate the truth of a belief if and only if the objective probability of the belief's being true, given that it's based on that ground, is sufficiently high (Alston 1988, 2005). Alston also takes the relevant notion of objective probability to be hypothetical relative frequency. Consider the

[8] Lam (2011) discusses a version of (Brier) (215-219). See also Dunn (2015) and Tang (2016).

[9] See, for example, Goldman (1999) (89).

[10] Incidentally, Alston (2005) argues that his version of indicator reliabilism is also a kind of process reliabilism—he thinks that 'reliability of process and reliability of indicator [. . . ] coincide' given certain plausible assumptions (137).

example above. What's the objective probability of our belief's being true given that it's based on our visual experience as of a cup on the table? On Alston's view, to answer this question, we should consider a reference class consisting of similar cases in which we form a similar belief that's based on a similar ground.[11] The frequency of such cases in which the relevant belief is true yields the value of the objective probability in question.

Given all of the above, it's natural to think that whether a credence is justified depends on whether it matches the corresponding objective probability. Exploiting this idea, here's yet another theory of justified credence:

**(Grounds)** For any $x \in [0, 1]$ and any $p$, S's credence of $x$ in $p$ is justified if and only if it's based on some ground $g$, where the objective probability of the credence having a true content given that it's based on $g$ equals (or approximates) $x$.[12]

Suppose for instance that S's credence of 0.8 in $p$ is justified. Then it's based on some ground $g$. Now consider the reference class consisting of cases similar to that in which S assigns a credence of 0.8 to $p$ based on $g$. According to (Grounds), the frequency of such cases in which S assigns a credence of 0.8 to a *true* proposition equals (or approximates) 80%.

As stated, (Grounds) deals only with cases in which the relevant grounds are non-doxastic in nature. But it can be extended to deal with cases in which the grounds are doxastic. (This will be relevant to the discussion in section 4.) Take a modified version of Black Box in which we'e certain that exactly 85% of the balls in the box are red. A proponent of (Grounds) might suggest that a credence of 0.85 in RED based on such certainty is justified if and only if (i) our certainty that 85% of the balls are red is justified and (ii) the objective probability of RED being true given that it's based on such justified certainty is 0.85.[13] (Whether the relevant

[11]There's an issue regarding what counts as a similar belief or ground. This issue is related to the generality problem—see Conee and Feldman (1998)–but we don't need to pursue it here. Again, my aim isn't to give a full-fledged defence of reliabilism.

[12]For discussion of versions of (Grounds), see Tang (2016) and Pettigrew (forthcoming).

[13]Process reliabilists make a somewhat similar move when dealing with inferential justification. Reasoning by modus ponens, we infer $q$ from our belief that $p$ and our belief that $q$ if $p$. The

state of certainty is justified will in turn depend on the ground on which it is based. If such a ground is doxastic in nature, we'll have to ask whether it is justified, and so on.)

## 3 Reliabilism and Imprecision

(Probability), (Calibration), (Brier), and (Grounds) are meant to account for precise credences. Can they be extended to deal with imprecise credences? Let's consider each theory in turn.

### 3.1 (Probability) and Imprecision

According to (Probability), a credence is justified just in case it's produced by a process that tends to produce a high proportion of credences that match the corresponding objective probabilities. One might suggest that an imprecise credence is justified just in case it's produced by a process that tends to produce a high proportion of credences, be they precise or imprecise, that match the relevant objective probabilities, be they precise or imprecise.[14]

To illustrate, suppose a process produces a precise credence of 0.7 in rain tomorrow. Then ideally, according to the suggestion, the objective probability of rain tomorrow should be 0.7 too. And if the process produces an imprecise credence of [0.6, 0.8] in rain tomorrow, then ideally, the objective probability of rain tomorrow should have the imprecise value [0.6, 0.8] too. The intuitive idea is that, while precise credences aim at precise objective probabilities, imprecise credences aim at imprecise objective probabilities.

However, the suggestion won't do. By its lights, if a process produces imprecise credences that don't match the corresponding objective probabilities, then so much

process of reasoning by modus ponens isn't reliable per se since whether it yields a high ratio of true beliefs depends on whether the input beliefs are true. But the process is *conditionally reliable*—it produces a high ratio of true beliefs given that all the belief inputs are true. Process reliabilists might thus hold that our belief that $q$ is justified if it is produced by a conditionally reliable process, and our beliefs that $p$ and that $q$ if $p$ are themselves justified.

[14]Fenton-Glynn (2019) argues that we should believe that imprecise objective chances exist (conditional on our adopting the Best System Analysis).

the worse for the process with respect to producing justified doxastic states. But this is a problem because imprecisers typically hold that we may have justified imprecise credences in propositions whose objective probabilities of being true *are* precise.

Consider Black Box. Suppose that the objective probability of RED—of picking a red ball from the box—is 85%, although we are not apprised of such information, and we know only that 80% to 90% of the balls are red. In this case, the objective probability of RED is precise. But according to Sturgeon, the corresponding credence ought to be imprecise. If he's right, then contrary to the current suggestion, a process may often produce justified imprecise credences in propositions whose objective probabilities are precise.

It also won't do to suggest that an imprecise credence is justified if and only if it's produced by a process that produces a high proportion of precise credences that match the corresponding objective probabilities. On this suggestion, whether an imprecise credence produced by a process is justified depends solely on the precise credences that the process produces. But imagine a process that produces precise credences that match the objective probabilities exactly, but whenever the relevant evidence is just slightly lacking in specificity, produces a maximally imprecise credence of [0, 1]. On the current suggestion, such a credence would be justified. But, presumably, even imprecisers won't want to say that it is.

### 3.2 (Calibration) and Imprecision

Does (Calibration) fare better? To deal with imprecise credences, one might suggest—naively—that a process is calibrated if and only if it satisfies two conditions: first, (i) for any $y \in [0, 1]$, $100y\%$ of the propositions in which it causes us to have a credence of $y$ are true, and second, (ii) for any $x, y \in [0, 1]$, [100x%, 100y%] of the propositions in which it causes us to have a credence of $[x, y]$ are true.

This suggestion can't work. Given condition (ii), a process that produces imprecise credences can't be calibrated. Consider all the relevant cases in which we assign a particular imprecise credence to a proposition. The proportion of cases in which we assign such a credence to a true proposition is *not* interval-valued. For example, suppose a credence-forming process causes us to assign a credence of [0.5, 0.7] to each of ten propositions. The proportion of true propositions is a precise value. Since, on the suggestion, a process that produces imprecise credences can't be calibrated, the doxastic states it produces won't be justified. But this runs contrary to the view that there are justified imprecise credences.

What if we modify condition (ii) as follows: for any $x, y \in [0, 1]$, the proportion of true propositions out of the propositions to which the process assigns an imprecise credence of $[x, y]$ falls within the interval $[100x\%, 100y\%]$? Satisfying this condition doesn't require that the proportion of true propositions be imprecise. Unfortunately, the modification faces another worry—it makes it far too easy for a process to count as being well calibrated and hence, far too easy for an imprecise credence to count as being justified.

To see why, suppose that a process assigns an imprecise credence of [0.4, 0.6] to each of ten propositions, exactly 50% of which are true. Given that '50%' falls within the interval [40%, 60%], the process counts as being calibrated by the lights of the current suggestion (at least with respect to those ten propositions). But by the same lights, the process wouldn't have been any less calibrated if it had instead produced any of the following credences, among infinitely many credences of other values, in each of those ten propositions: [0, 0.5], [0.49, 0.51], [0.1, 0.9], or even [0.5, 0.5].

Also, the current suggestion rewards imprecision rather indiscriminately: instead of assigning a credence of [0.4, 0.6] to each of a bunch of propositions, it would be better, given the suggestion, to assign a credence of, say, [0.1, 0.9] to each of them—never mind what evidence we have. For given the suggestion, increasing the imprecision of the credences produced by a process will only make it easier for the process to be calibrated. In the extreme, a process that produces only maximally imprecise credences is guaranteed to be perfectly calibrated (since the relevant proportion of true propositions is guaranteed to fall within the interval [0%, 100%]). But we shouldn't fetishise imprecision. Presumably even imprecisers will hold that sometimes our credences ought to be more rather than less precise.

### 3.3 (Brier) and Imprecision

Both (Probability) and (Calibration) face difficulty accounting for imprecision. How well does (Brier) fare? The Brier score allows us to calculate the accuracy of precise credences. But what about imprecise ones? I'll first consider some suggestions on how to extend (Brier) to account for justified imprecise credences and show why they are problematic. I'll then appeal to work by Schoenfield (2017) that helps show why, in general, any attempted extension will face difficulty.

Call any score that is meant to deal with imprecise credences and that is inspired by the Brier score a *Brier* score*. Seidenfeld et al. (2012) suggest that the Brier* score of an imprecise credence of $[x, y]$ in $p$ be given by $(1-y)^2$ if $p$ is true and $x^2$ if $p$ is false (1252). (When we have a precise credence of $x$, its Brier* score will just be its Brier score.) Given the suggestion, one might hold that, for any $x, y \in [0, 1]$ and any $p$, a credence of $[x, y]$ in $p$ is justified if and only if it is produced by a process that tends to produce credences that have a low average Brier* score.

The suggestion runs into difficulty. A version of (Brier) that adopts such a scoring method is vulnerable to a *systematic* bias against the assignment of precise credences. For we'll *always* get a better score if, instead of assigning a sharp credence of $z$ to $p$, we assign an imprecise credence of $[x, y]$ to it, where $x < z < y$. To illustrate, whether $p$ is true or false, we'll get a better score if, instead of assigning a sharp credence of 0.6 to $p$, we assign an imprecise credence of [0.5, 0.61] to it. For if $p$ is true, then assigning it a credence of 0.6 yields a score of 0.16 whereas assigning it a credence of [0.5, 0.61] yields a score of 0.1521, but if $p$ is false, then assigning it a credence of 0.6 yields a score of 0.36 whereas assigning it a credence of [0.5, 0.61] yields a score of 0.25.

More generally, on the current suggestion, decreasing the lower value and increasing the upper value of an imprecise credence guarantees a decrease in its Brier* score. For example, suppose that we assign a credence of [0.1, 0.8] instead of a credence of [0.5, 0.61] to $p$. Then if $p$ is true, the associated Brier* score is 0.04, but if $p$ is false, the associated Brier* score is 0.01. In either case, the Brier* score of a credence of [0.1, 0.8] in $p$ is higher than that of a credence of [0.5, 0.61] in $p$. And in the extreme case, a process that generates a maximally imprecise credence in any proposition whatsoever will result in an average Brier* score of 0. But it's implausible that we should *always* penalise precision in favour of imprecision. Even if some situations require less precision, some require more. For instance, if we know for certain that the chance of $p$ is 0.8 (and have no other relevant evidence concerning the truth of $p$), it seems that we ought to assign a precise credence of 0.8 to $p$.[15]

It also won't help to hold that the Brier* score of an imprecise credence of $[x, y]$ in $p$ is just the Brier score of a credence of $x$ in $p$, or the Brier score of a credence of $y$ in $p$, or the Brier score of a credence in $p$ that is the average of $x$ and $y$.[16] For then, there would be no epistemic difference between assigning a credence of $[x, y]$ to $p$ and assigning the corresponding (precise) credence to $p$. But imprecisers would presumably hold that there is a difference.

Can we avoid the problems above? Perhaps, to calculate the Brier* score of one's credence, we should use the scoring rule put forward by Seidenfeld et al. (2012) but also add a certain value $w$ to the score, where $w$ measures the relevant degree of imprecision. The exact size of $w$ will depend on how much precision is valued, but by imposing a penalty on imprecision, we at least avoid the result that a maximally imprecise credence has a Brier* score of 0.

Further, depending exactly on how $w$ is computed, penalising imprecision may

[15]A similar problem arises if we take the Brier* score of an imprecise credence of $[x, y]$ in $p$ to be given by $(1 - x)^2$ if $p$ is true and $y^2$ if $p$ is false. Whereas the scoring method proposed by Seidenfeld et al. (2012) systematically favours imprecision, this scoring method systematically favours precision. No matter what score we get by assigning an imprecise credence of $[x, y]$ to $p$, we are guaranteed to get a lower score if we assign a precise credence of $z$ to $p$, where $x < z < y$.

[16]For example, as Mayo-Wilson and Wheeler (2016) point out, a credence of 0.5 and a credence of [0, 1] will be equally accurate if we score an imprecise credence such as [0, 1] by its midpoint (67).

also help us avoid a systematic bias against the assignment of precise credences. For example, suppose that for any imprecise credence of $[x, y]$ in $p$, $w$ is calculated by taking the square of the difference between $x$ and $y$. Then if $p$ is true, a credence of [0.5, 0.61] in $p$ will yield a score of 0.1642 whereas a credence of 0.6 in $p$ will yield a score of 0.16. And if $p$ is false, the former will yield a score of 0.2621, whereas the latter will yield a score of 0.36. So if $p$ is true, a credence of [0.5, 0.61] in $p$ will yield a higher Brier* score than a credence of 0.6 in $p$, but if $p$ is false, the former will yield a lower Brier* score than the latter.

Unfortunately, this suggestion runs into difficulty too. Consider an assignment of an imprecise credence of [0.1, 0.9] to $p$. On the current suggestion, whether $p$ is true or false, such an imprecise credence will be assigned a Brier* score of $w + 0.01$, where the value of $w$ will depend on how much we value precision. Now, the Brier* score of a precise credence of 0.5 in $p$ is 0.25 whether $p$ is true or false. Question: should $w$ be equal to 0.24, less than 0.24, or greater than 0.24? Every option leads to a problem.

If $w$ equals 0.24, there would be *no* difference, epistemically speaking, between assigning a precise credence of 0.5 to $p$ and an imprecise credence of [0.1, 0.9] to $p$ (since both would yield the same score). If $w$ is less than 0.24, then it's *always* better to assign an imprecise credence of [0.1, 0.9] rather than a precise credence of 0.5 to $p$. And if $w$ is greater than 0.24, it's *always* better to assign a precise credence of 0.5 rather than an imprecise credence of [0.1, 0.9] to $p$. But every option goes against the spirit of rooting for imprecise credences. For though imprecisers hold that some situations require an imprecise credence of [0.1, 0.9] in $p$, they'll presumably grant that some situations require a precise credence of 0.5 in $p$—it all depends on one's evidence regarding $p$. The suggestion under consideration is problematic because it doesn't allow such a possibility.[17]

The difficulty we face trying to extend (Brier) to deal with imprecise credences can be generalised. (Brier) is inspired by accuracy-first epistemology, according to

[17]Cf. Schoenfield (2017) and Levinstein (2019).

which the fundamental epistemic value is accuracy (Pettigrew 2016, 8). Accuracy-first epistemologists, as Schoenfield (2017) puts it, hold that '*all* epistemic norms are rooted fundamentally in an agent's rational pursuit of accuracy' and that '[t]he accuracy-centered epistemologist's project involves showing how rational requirements can be derived from accuracy-based considerations' (669). But Schoenfield has argued that accuracy-first epistemology is incompatible with the position that some situations require imprecise credences.[18]

For full details of the arguments, see Schoenfield (2017). But two results are relevant to our discussion of (Brier). Informally, the first result says that for any imprecise credence function defined over the partition $\{p, \neg p\}$, and for any acceptable numerical accuracy measure for imprecise credence functions (that satisfy certain plausible constraints), there will be a precise credence function that is no less accurate than the imprecise credence function in question, no matter how the world turns out to be (673). This result is a problem for attempts to extend (Brier) to accommodate imprecise credences. It's also a problem for versions of (Brier) that, instead of appealing to Brier scoring, appeal to any other acceptable numerical

[18]For related arguments, see Lindley (1982), Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2016), Berger and Das (forthcoming), and Levinstein (2019). Schoenfield also argues that accuracy-first epistemology is incompatible with the position that imprecise credences are sometimes *permitted*. Konek (2019) tries to make room for imprecision by relating it to how much an agent values closeness to the truth and disvalues distance from error. In particular, he rejects an assumption (that of *admissibility*, according to which probabilistic belief states are non-dominated) that is crucial to Schoenfield's argument for the incompatibility of accuracy-first epistemology with imprecision. On Konek's view, '[w]e see different lower and upper probabilities as appropriate responses to the same evidence not because we disagree about the strength of the evidence [...] but rather because we take different attitudes toward the comparative importance of avoiding error and pinning down the truth, and different types of lower/upper probabilities (intervals) do a better job at one or the other' (Konek 2019, 259). But as Carr (2016) notes, Konek's 'view does not take into account the character of an agent's evidence and the extent to which it is informative about objective chances'—thus ignoring a major motivation for imprecision (71). For example, with respect to Black Box, our relevant evidence is that exactly 80% to 90% of the balls in the box are red. On Sturgeon's view, any credence other than an imprecise credence of [0.8, 0.9] in our picking a red ball at random is unjustified. But on Konek's view, different imprecise attitudes may be appropriate depending on how much we value avoiding error and value pinning down the truth. Now, ultimately, Konek might be right and Sturgeon might be wrong. But Schoenfield is explicit that her target is what she calls the *Standard Imprecise View*, according to which 'if the only evidence you have concerning whether P is that the objective chance function for {P, ∼P} is in the set of probability functions $S$, then your evidence requires you to adopt the doxastic attitude represented by $S$.' (668). Further, as mentioned at the beginning of this paper, I'm assuming that Sturgeon's intuition about Black Box is correct. This paper explores whether reliabilism can be reconciled with imprecision on such an assumption.

accuracy measure.

To see why, suppose that a process produces imprecise credences in various propositions (and their negations), and the average accuracy of the credences produced by the process has a certain value $X$. Given the result above, this average accuracy would have been no less than $X$—no matter how the world turned out to be—if, instead of producing the imprecise credences in question, the process had produced the relevant precise credences. Sticking to the spirit of (Brier), which ties the justifiedness of the credences produced by a process with their accuracy, this means that if the process had produced the relevant precise credences, such credences would count as being no less justified than the corresponding imprecise credences. This runs counter to the position that some situations require imprecise credences.

In response, one might suggest that we have been barking up the wrong tree—that we should not have been scoring the accuracy of imprecise credences with precise numbers, but with, say, sets of numbers (Schoenfield 2017, 675). But here enters the second result. In Schoenfield's words, it says that 'no matter what sort of object we use to represent accuracy, the credence of 0.5 in each cell of a two-cell partition is no less accurate than any imprecise state that assigns to each cell in the partition an interval-valued credence centered at 0.5'—and this is supposed to hold regardless of how the world turns out to be (675). Given the second result, the suggestion won't help.

Here's why. Suppose that a process produces imprecise credences in various propositions $p$, $q$, $r$, ... that are centred at 0.5, say, credences represented by the intervals [0.1, 0.9], [0.2, 0.8], [0.45, 0.55], and so on. Given the preceding result, the average accuracy of the credences produced by the process would have been no worse—no matter how the world turned out to be—if, instead of producing the imprecise credences in question, the process had produced a credence of 0.5 in the relevant propositions. Sticking to the spirit of (Brier), which ties the justifiedness of the credences produced by a process with their accuracy, this means that if the process had produced a precise credence of 0.5 in each of $p$, $q$, $r$, ..., such precise credences would count as being no less justified than the relevant imprecise credences. This runs counter to the position that some situations require imprecise credences rather than a precise credence of 0.5.

### 3.4 (Grounds) and Imprecision

Let's now turn to (Grounds). One might suggest that we extend (Grounds) as follows: for any $x, y \in [0, 1]$ and any $p$, S's credence of $[x, y]$ in $p$ is justified if and only if it is based on some ground $g$, where the objective probability of the credence having a true content given that it's based on $g$ equals $[x, y]$. But the suggestion faces a problem similar to that faced by (Calibration) if, following Alston, we take 'objective probability' to refer to hypothetical relative frequency. Consider cases similar to that in which S assigns a credence of $[x, y]$ to $p$ based on ground $g$ (for some $x$, $y$, $p$, and $g$). The frequency of such cases in which S assigns a credence of $[x, y]$ to a true proposition is a sharp value. So the current suggestion has the implausible result that no imprecise credence is justified. For there would never be a match between an imprecise credence and the corresponding relative frequency.

It won't help to take the relevant notion of objective probability to be some sort of single-case objective chance. Consider Black Box once again. Suppose that, unbeknown to us, the objective chance of RED is 0.85. And suppose that we form an imprecise credence of [0.8, 0.9] in RED based on our certainty that the chance of RED is between 80% and 90%. The chance of RED, given that our imprecise credence is based on the relevant ground, is still 0.85. Our basing our imprecise credence on our certainty that the chance of RED is between 80% and 90% does not make the objective chance of RED imprecise.

Let me elaborate, but with the help of a slightly different example. Suppose that the objective chance of a radioactive atom decaying in the next minute is precisely 0.5. But suppose we know only that its chance of decaying is between 0.4 and 0.6. Note that the chance of decay is independent of whether we know its exact value—

our knowing only that the chance is between 0.4 and 0.6 does not affect the chance of decay. In other words, the chance of decay conditional on our knowing only that such a chance is between 0.4 and 0.6 is exactly 0.5 and not some imprecise value. Similarly, our being certain that the chance of RED is between 80% and 90% does not affect the objective chance of RED. If it is precise to begin with, it remains precise conditional on the ground in question.[19]

## 4 Reliabilism and Indeterminate Credences

We've seen four different ways for reliabilists to account for precise credences. We've also looked at several attempts to extend such accounts to deal with imprecise credences. But all the attempts face problems. This gives us reason to think that reliabilism is incompatible with imprecision—at least if we take imprecise credences to be determinate first-order attitudes. As mentioned, such incompatibility is significant. First, it's not initially obvious that reliabilists should reject imprecision. Second, contrary to a natural thought, imprecisers can't be neutral as to which theory of doxastic justification is correct.

But what about the intuition that in Black Box, the relevant doxastic state should be associated with an interval instead of a precise number? Must reliabilists reject it entirely? I don't think so, and I'd like to conclude by exploring a way in which reliabilists might accommodate the intuition.

Recall Levinstein's interpretation of imprecise credences. According to him,

[19]One might suggest that a proponent of (Grounds) should appeal to *epistemic* probability instead of objective probability. Suppose that there are epistemic probabilities and that such probabilities may be imprecise. Then it's natural to suggest that your credences ought to be imprecise when the relevant epistemic probabilities are imprecise. (Comesaña (2018) endorses a theory similar to (Grounds) but maintains that the relevant notion of probability should be epistemic probability. He does not have imprecise credences in mind; instead he thinks that such a theory will avoid certain other problems faced by versions of (Grounds) that appeal to objective probability.)

However, the suggestion incurs a price that might be too hefty for most reliabilists. Reliabilists, following Goldman (1979), typically want to cash out justification in non-epistemic terms. But *epistemic probability* is an epistemic notion; further, it's notoriously difficult to cash it out in non-epistemic terms. As Williamson (2000) puts it, 'Carnap's programme of inductive logic is moribund', and any attempt to spell out epistemic probability in purely syntactic terms is doomed, since the 'difference between green and grue is not a formal one' (Williamson 2000, 211).

> when we say that an agent has imprecise credence [.2, .3] towards some proposition $X$, we don't mean that her credence is literally the interval [.2, .3]. Instead, there's no fact of the matter as to whether her credence is really .22, .29 or any other element of [.2, .3]. (742)

Levinstein thinks that such an interpretation escapes Schoenfield's arguments for the incompatibility of accuracy-first epistemology and imprecision. And while Schoenfield stops short of endorsing the interpretation, she grants that an 'alternative to the imprecise credence model is to think that it can sometimes be *indeterminate* what doxastic attitude an agent takes (or should take) towards a proposition' (680; Schoenfield's emphasis).[20]

I don't wish to be caught up in the question whether imprecise credences should ultimately be understood as indeterminate credences. Even if you don't like such an interpretation, it's still worth exploring whether indeterminate credences are, in their own right, compatible with reliabilism.[21] Suppose we don't think that there are imprecise credences understood as determinate first-order attitudes. Does it follow immediately that, with respect to Black Box, we should assign a precise credence to RED?

One might think not. As Schoenfield puts it, 'denying that sets of credence functions represent genuine doxastic alternatives doesn't *entail* that agents have, or ought to have, precise credences in every proposition' (680; Schoenfield's emphasis). But suppose that's right. Then there's room for holding that we ought to have an indeterminate credence in RED. More carefully, there's room for holding that we ought to be in a state in which it's indeterminate whether our credence in RED is 0.8, or 0.85, or any other precise value in the interval [0.8, 0.9]—as opposed to, say, a state in which the indeterminacy ranges over some other interval or one in which we determinately assign a precise credence to RED. This helps accommodate the

[20] Also, see Rinard (2015), who provides a decision theory for imprecise credences understood as indeterminate credences.

[21] For brevity's sake, talk of one having an indeterminate credence of $[x, y]$ in $p$ should be understood as talk of it being indeterminate what precise credence one has, where such indeterminacy is represented by the interval $[x, y]$.

intuition that in Black Box, the relevant doxastic state should be associated with the interval [0.8, 0.9] instead of a precise number or some other interval.

Question: can reliabilists account for the kind of indeterminacy mentioned above? I think so, and to show this, it suffices to show that at least one of the theories discussed in the previous section is up to the task. I'll focus on (Grounds), which I think can accommodate indeterminacy naturally.[22]

Earlier we considered an extension of (Grounds) according to which an imprecise credence is justified only if the corresponding objective probability, understood as hypothetical relative frequency, is similarly imprecise. But recall the worry that such a match is unobtainable since the corresponding frequency relative to a certain reference class is always precise. Can we avoid this problem if we focus on indeterminate credences?

Here's another extension of (Grounds): S's indeterminate credence of $[x, y]$ in $p$ is justified if and only if it's based on some ground $g$ and the objective probability of $p$'s being true, given that S's credence is based on $g$, is indeterminate over the interval $[x, y]$.

This extension—call it *(Grounds*)*—is different from that considered previously. Suppose that S has an indeterminate credence of $[0.4, 0.6]$ in $p$ that's based on $g$. (In other words, suppose that, in response to $g$, S enters a state in which, for any $x \in [0.4, 0.6]$, it's indeterminate whether S's precise credence is $x$.) Suppose further that we've picked out the relevant reference class that consists of cases similar to that in which S forms a credence of $[0.4, 0.6]$ in $p$ based on $g$. No doubt, relative to this reference class, the frequency of S forming a credence in a true proposition is

[22]Not all reliabilist theories of credences can accommodate indeterminacy easily. Recall that, according to (Probability), a justified credence in $p$ comes from a process that produces a high proportion of credences that match the corresponding objective probabilities. Now suppose that it's indeterminate what (precise) credence we have in RED. One might suggest that, ideally, by the lights of (Probability), it should be indeterminate what the objective probability of RED is. But this doesn't seem right. After all, the proportion of red balls in the box before us is a determinate value—it's just that we don't know what it is. So, though the exact objective probability of RED is unknown, it has a determinate value too. In general, a proponent of indeterminate credences might maintain there are many cases like Black Box in which the objective probability of the relevant proposition is determinate, but it's false that one ought to have a determinate (precise) credence in that proposition. (I'll leave it to the reader to think about whether (Calibration) and (Brier) can accommodate indeterminate credences.)

a sharp value. Nonetheless, *it may sometimes be indeterminate what the reference class is.* It might be that given a certain reference class $C_1$, the relevant frequency is 0.4, that given another reference class $C_2$, the relevant frequency is 0.6, and that given yet another reference class $C_3$, the relevant frequency is some value between 0.4 and 0.6. But suppose it's indeterminate whether the relevant reference class is $C_1$, $C_2$, or $C_3$. Then it's indeterminate whether the relevant objective probability is 0.4, 0.6, or some value between 0.4 and 0.6.

Why might it sometimes be indeterminate what the reference class is? Consider a modified version of Black Box. Suppose we're certain that exactly 85% of the balls are red, and we form a credence of 0.85 in RED based on such a ground. Is our credence in RED justified?

Recall our discussion of (Grounds) in section 2.4. Proponents of (Grounds) may hold that our credence in RED is justified if and only if (i) our certainty that 85% of the balls are red is justified and (ii) the objective probability of RED being true given that it's based on such justified certainty equals 0.85. Suppose that our certainty that 85% of the balls are red is indeed justified. Then it remains to see if the relevant objective probability is 0.85.

To do this, the reference class to consider should consist of cases in which we form a credence of 0.85 in RED based on justified certainty that 85% of the balls in the box before us are red. Now, given (Grounds), our certainty that $p$ is justified only if $p$. So the relevant reference class should consist of cases in which exactly 85% of the balls in the box before us are red. But then the frequency with which we draw a red ball relative to such a reference class will tend to 85%. There's a match between our credence of 0.85 in RED and the corresponding objective probability.

But now, suppose that we're justifiedly certain that 80% to 90% of the balls in the box are red (and we're apprised of no other relevant information). Suppose also that, based on such a ground, we form an indeterminate credence of [0.8, 0.9] in RED.[23] Plausibly, the relevant reference class should consist of similar cases in

[23] The expression 'forming an indeterminate credence of [0.8, 0.9] in $p$' should be interpreted in the way mentioned in footnote 21.

which we form an indeterminate credence of [0.8, 0.9] in our randomly drawing a red ball based on our being justifiedly certain that exactly 80% to 90% of the balls in the box before us are red. Further, such cases will be ones in which exactly 80% to 90% of the balls in the box before us are red. What is the frequency of such cases in which we draw a red ball?

Our answer can't have the form of a precise value. We can't say that the frequency in question is 80% or that it is 85%, for instance. To say that would be to presume that there is a determinate reference class we can pick out consisting of (a) cases in which exactly 80% of the balls in the box are red, or one consisting of (b) cases in which exactly 85% of the balls in the box are red. But we have no principled reason to pick one over the other. For the reference class is determined in part by the ground on which a subject's credence is based, and the relevant ground—our being justifiedly certain that 80% to 90% of the balls are red—does not in itself favour (a) over (b) or vice versa.

However, given the ground in question, for any $x \in [0, 0.8) \cup (0.9, 1]$, we can determinately rule out the reference class containing cases in which $100x\%$ of the balls are red. So we can say that the objective probability of RED (given the relevant ground) is indeterminate over the interval [0.8, 0.9]—instead of, say, the interval [0.5, 0.8]. Our indeterminate credence of [0.8, 0.9] in RED is then justified by the lights of (Grounds*), since the indeterminacy in our credal state is matched by a similar indeterminacy in the relevant objective probability.

Might there be a determinate reference class we can pick out—namely, one containing cases in which the percentage of red balls equals the percentage of red balls in the *actual* situation? Suppose that as a matter of fact, 85% of the balls in the box before us are red. But suppose that we're not apprised of this fact and know only that 80% to 90% of the balls are red. Nonetheless, on the current suggestion, the relevant reference class is just one consisting of cases in which 85% of the balls before us are red.

But this line of thought can't be right, or (Grounds) would be in trouble even with respect to precise credences. Suppose we're justifiedly certain that the percentage of red balls in the box before us is either 85% or 20%, but we don't know which and we don't know anything else relevant. Tom, whom we are justifiedly certain is 99% reliable, comes along. He informs us that the percentage of red balls is 20% and not 85%. It seems that we would be justified in having a low credence (0.2065, to be precise) in our picking a red ball. But suppose that this is a rare occasion on which Tom is wrong, and in fact, 85% of the balls are red. The suggestion under consideration yields the implausible result that a low credence would be unjustified, and that we should assign a credence of 0.85 to our picking a red ball. For given the suggestion, the reference class should just consist of cases in which 85% of the balls are red. In this case, the relevant objective probability of picking a red ball would equal 0.85.

In sum: the relevant reference class is determined by our ground and not by the actual percentage of red balls in the box before us. When our ground consists only of our being justifiedly certain that exactly 80% to 90% of the balls are red, there is just no fact of the matter whether the relevant reference class contains cases in which 80% or 85% or 90% of the balls are red. So it's indeterminate what the objective probability of RED (given the relevant ground) is, and such indeterminacy can be reflected by its being indeterminate what credence we assign or ought to assign to RED. Reliabilists can thus accommodate the intuition that the relevant doxastic state in Black Box should be associated with an interval rather than a precise number.[24]

[24] I thank Ben Blumson, Michael Pelczar, and Jens Christian Bjerring for comments on an early version of this paper. I also thank the participants of an epistemology workshop at the National University of Singapore, the audience at the 2014 International Conference on Epistemology and Cognitive Science at Xiamen University, the participants of a reading group organised by Richard Pettigrew at the University of Bristol, and the audience at a King's College London seminar organised by Julien Dutant. Last but not least, I thank an anonymous referee of this journal for their helpful feedback.

**References**

Alston, W. P. (1988). An internalist externalism. *Synthese 74*, 265–83.

Alston, W. P. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation.* US: Cornell University Press.

Berger, D. and N. Das (forthcoming). Accuracy and credal imprecision. *Nous*.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review 78*, 1–3.

Carr, J. (2016). Chancy accuracy and imprecise credence. *Philosophical Perspectives 29*, 67–81.

Carr, J. (ms.). Imprecise evidence without imprecise credences.

Comesaña, J. (2018). Whither evidentialist reliabilism? In K. McCain (Ed.), *Believing in Accordance with the Evidence: New Essays on Evidentialism*, Volume 398 of *Studies in Epistemology, Logic, Methodology, and Philosophy of Science.* Springer International Publishing.

Conee, E. and R. Feldman (1998). The generality problem for reliabilism. *Philosophical Studies 89*, 1–29.

Dunn, J. (2015). Reliability for degrees of belief. *Philosophical Studies 172*, 1929–1952.

Elga, A. (2010). Subjective probabilities should be sharp. *Philosophers' Imprint 10*, 1–11.

Fenton-Glynn (2019). Imprecise chance and the best system analysis. *Philosophers' Imprint 19*, 1–44.

Goldman, A. I. (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and Knowledge: New Studies in Epistemology*, pp. 1–23. Dordrecht: D. Reidel Publishing Co.

Goldman, A. I. (1986). *Epistemology and Cognition.* Cambridge, MA: Harvard University Press.

Goldman, A. I. (1999). *Knowledge in a Social World.* New York: Oxford University Press.

Goldman, A. I. (2012). *Reliabilism and Contemporary Epistemology: Essays.* New York: Oxford University Press.

Hájek, A. (MS). A puzzle about partial belief.

Hájek, A. and M. Smithson (2012). Rationality and indeterminate probabilities. *Synthese 187*, 33–48.

Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science 65*, 575–603.

Joyce, J. M. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives 24*, 281–323.

Konek, J. (2019). Ip scoring rules: Foundations and applications. *Proceedings of Machine Learning Research 103*, 256–264.

Konek, J. (forthcoming). Epistemic conservativity and imprecise credence. *Philosophy and Phenomenological Research.*

Lam, B. (2011). On the rationality of belief-invariance in light of peer disagreement. *The Philosophical Review 120*, 207–245.

Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy 71*, 391–418.

Levi, I. (1985). Imprecision and indeterminacy in probability judgment. *Philosophy of Science 52*, 390–409.

Levinstein, B. A. (2019). Imprecise epistemic values and imprecise credences. *Australasian Journal of Philosophy 97*, 741–60.

Lindley, D. V. (1982). Scoring rules and the inevitability of probability. *International Statistical Review 50*, 1–11.

Mayo-Wilson, C. and G. Wheeler (2016). Scoring imprecise credences: A mildly immodest proposal. *Philosophy and Phenomenological Research 153*, 55–78.

Pappas, G. S. (Ed.) (1979). *Justification and Knowledge: New Studies in Epistemology.* Dordrecht: D. Reidel Publishing Co.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence.* Oxford: Oxford University Press.

Pettigrew, R. (forthcoming). What is justified credence? *Episteme*.

Rinard, S. (2015). A decision theory for imprecise probabilities. *Philosophers' Imprint 15*, 1–16.

Schoenfield, M. (2017). The accuracy and rationality of imprecise credences. *Nous 51*, 667–85.

Seidenfeld, T., M. J. Schervish, and J. B. Kadane (2012). Forecasting with imprecise probabilities. *International Journal of Approximate Reasoning 53*, 1248–1261.

Sturgeon, S. (2010). Confidence and coarse-grained attitudes. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, Volume 3, pp. 126–49. Oxford University Press.

Tang, W. H. (2016). Reliability theories of justified credence. *Mind 125*, 63–94.

van Fraassen, B. C. (1983). Calibration: A frequency justification for personal probability. In R. S. Cohen and R. Laudan (Eds.), *Physics, Philosophy and Psychoanalysis*, Volume 76 of *Boston Studies in the Philosophy of Science*, pp. 295–319. In Cohen and Laudan 1983.

van Fraassen, B. C. (1984). Belief and the will. *The Journal of Philosophy 81*, 235–256.

van Fraassen, B. C. (1990). Figures in a probabilistic landscape. In J. M. Dunn and A. Gupta (Eds.), *Truth or Consequences: Essays in Honor of Nuel Belnap*, pp. 345–356. Dordrecht: Kluwer Academic Publishers.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities, Monographs on Statistics and Applied Probability.* London: Chapman and Hall.

White, R. (2010). Evidential symmetry and mushy credence. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, Volume 3. Oxford University Press.

Williamson, T. (2000). *Knowledge and Its Limits.* Oxford University Press.