Rational Agnosticism and Degrees of Belief *

Jane Friedman

Abstract and Keywords

There has been much discussion about whether traditional epistemology's doxastic attitudes are reducible to degrees of belief. In this paper it is argued that what we may call the Straightforward Reduction—the reduction of all three of believing p, disbelieving p, and suspending judgment about p, ¬p to precise degrees of belief for p, ¬p that ought to obey the standard axioms of the probability calculus—cannot succeed. By focusing on suspension of judgment (agnosticism) rather than belief, we can see why the Straightforward Reduction is bound to fail. It is argued that, in general, suspending about p is not just a matter of having some specified standard credence for p, and in the end some ways to extend the arguments are suggested that will put pressure on other credence-theoretic accounts of belief and suspension of judgment as well.

Keywords: Belief, credence, degree of belief, subjective probability, suspension of judgment, agnosticism, Lockean Thesis, ignorance, doxastic attitudes

1. Preliminaries

"Traditional" epistemology presents us with a quite spare doxastic taxonomy. This taxonomy is typically thought to include just belief, disbelief, and suspension of judgment. And given that in this context disbelieving p is thought identical to believing $\neg p$, traditionalists present us with just two attitudes: belief and suspension of judgment.

"Formal" epistemology presents us with an expansive doxastic taxonomy: a subject faces a continuum-many different doxastic options with respect to p. These doxastic states are thought to be something like states of confidence in the truth of the relevant propositions. These states of confidence come in degrees and those degrees of belief (credences) are taken to bear some intimate relation to the probability calculus.

A number of interesting questions arise about the relationship between these two doxastic taxonomies. An obviously pressing one is which (if any) are accurate: do both of these doxastic taxonomies accurately describe our doxastic lives, or does just one (or does neither)? If only one gives a correct description, which one, and what should we say about the other? Many have thought that only one of these taxonomies is accurate and that the attitudes described in the other should be identical or reducible to the attitudes described in the correct taxonomy. Those keen on such a reduction typically take the formalist's taxonomy to correctly describe our doxastic lives and aim to reduce the traditionalist's doxastic attitudes to degrees of belief. The formalist's taxonomy can seem much richer than the traditionalist's, and as such the former is often thought capable of simply subsuming the latter.

Let's define the *Straightforward Reduction* as follows. The Straightforward Reduction is a reduction of the traditionalist's doxastic attitudes to the (p.58) formalist's degrees of belief that says that believing p, disbelieving p, and suspending about p are just matters of having (some specified) *standard credences* for p.[1] What is a standard credence? Credences will count as standard for our purposes if the following two assumptions hold:

1. A subject's total doxastic standing is represented with a single credence function.
2. The standard (Kolmogorov) axioms of the probability calculus are normative for credence functions.

Here is the basic picture of standard credences with which I will be working. Take all the propositions for which a subject S has some degree of confidence or other. S's confidence in these propositions is

modeled or represented with a single real-valued function that takes these propositions as arguments. This is S's credence function ($C_s(\cdot)$). The numbers in the range of this function are S's degrees of belief or levels of confidence in the relevant propositions.[2] This function represents all of S's degrees of belief at a given time—what I am calling her total doxastic standing. The standard axioms of the probability calculus are normative for this credence function.[3] A subject with a credence function that fails to be a probability function is (at least to some extent) irrational. Credence functions aren't always probability functions, but they ought to be probability functions.[4]

In this paper I want to argue that the Straightforward Reduction simply cannot succeed. Most discussions of the traditional-to-formal reduction (and the relationship between the two taxonomies in general) have focused on traditional belief. Instead, I want to focus on the traditionalist's other attitude: suspended judgment (agnosticism).[5] I want to argue that attempting to reduce suspension about p to having a standard credence for p (in some range) (p.59) guarantees the failure of the Straightforward Reduction, and shows why suspension about p is not (just) a matter of having a standard credence for p. I will also say a bit about how we might extend some of the arguments in this paper to make more trouble for common credence-theoretic accounts of belief and other credence–theoretic accounts of suspension. In general, thinking about the p-credences of a rational p-agnostic can significantly impact one's thinking about the relationship between the two doxastic taxonomies.[6]

My focus then is on the traditionalist's second attitude, suspension of judgment. But should we even think of suspension as an attitude? I think that we should.[7] Suspended judgment is, or at least involves, a proper attitudinal doxastic commitment. Agnosticism about p is not merely failing to believe both of p and $\neg p$. We have the property of neither believing nor disbelieving all sorts of propositions about which we are not agnostic: propositions we cannot grasp, or those we can but have never contemplated or had in mind in any way. Suspension requires some sort of decision about or commitment with respect to the truth of p; it isn't a state that we are in in virtue of being opinionless, rather it is a state of opinion. It is in this sense that suspension is, or at least involves, a proper doxastic commitment about the truth of p on the part of the subject. The most natural way of understanding this commitment is as an attitude. What sort of attitude? A subject who suspends is effectively neutral or undecided

about whether p is true. Her attitude then is one that represents or expresses or simply is her neutrality or indecision with respect to the relevant content.[8]

At first glance it might seem as though we have a good chance of finding this "indecision-representing" attitude in the formalist's expansive taxonomy. After all, according to that taxonomy there are many intermediate doxastic attitudes: states of confidence that fall between a subject's being absolutely confident in the truth of p and her being absolutely confident in the truth of $\neg p$. In order to show that the Straightforward Reduction cannot succeed and that suspending about p is not (just) a matter of having a standard credence for p, I will rely on a very basic principle of rational suspension of judgment. I will argue that given this norm for suspension, if we try to say that suspending about p is just a matter of having a standard credence for p, we will have to say that it's a matter of having any standard credence for p at all; that is, we will have to say that one suspends about p iff $C_s(p) \varepsilon [0,1]$.[9]

(p.60) The norm for rational suspension that I want to focus on can be called the *absence of evidence* norm. This norm says roughly that in the absence of evidence for or against an ordinary contingent proposition p, it is epistemically permissible to suspend judgment about p. Let me clarify and defend this norm.

I want to rely on a norm for suspension that is relatively uncontroversial: that suspension about p is epistemically permissible if you have no evidence relevant to whether p is true or false. But in order for this norm to be largely uncontroversial we need first to limit the class of propositions it applies to. It would be fairly controversial to claim that suspension about (say) '2 + 2 = 4' is epistemically permissible absent evidence relevant to whether it is true or false. There are other sorts of propositions for which the absence of evidence norm might be more controversial as well. Some immediate perceptual propositions, or introspective ones might not be good candidates for falling under the norm. And some other sorts of contingent propositions might be able to escape as well: Kripke-style superficially contingent a priori propositions, and perhaps even some deeply contingent ones.[10] Nonetheless, it is easy to see that for a wide class of propositions the absence of evidence norm is fairly uncontroversial. The propositions to focus on are what I'll call *ordinary contingent propositions*. These propositions are "deeply" contingent for the relevant subjects—subjects will have no semantic

guarantee of their truth—and they simply describe mundane (possible) facts about the physical world.

Here is one helpful way to think about the absence of relevant evidence. Cases in which S lacks evidence relevant to p are cases in which S has no evidence either for or against p. In these sorts of cases, while S may have information (e.g. semantic or conceptual) about p, none of that information favours one of p or $\neg p$ over the other. What's important here is not that S's total evidence fail to favour one of p or $\neg p$ over the other. This can happen when that total body of evidence has some bits that favour p over $\neg p$ and other bits that favour $\neg p$ over p, and these bits are evenly weighted. These sorts of evidential circumstances also look like circumstances in which S is permitted to suspend judgment, but they aren't circumstances in which S has no relevant evidence. This is precisely because some bits of S's total evidence do favor one of p, $\neg p$ over the other. Subjects who lack evidence for or against p possess no information that supports one of p or $\neg p$ over the other.

This idea is sometimes fleshed out in probabilistic terms. Evidence that is relevant to p is sometimes thought to be the sort of evidence that induces a change in a rational subject's degree of belief for p.[11] One way to think of a subject who has no evidence for or against p is as a rational subject whose credence for p has not been moved at all by the evidence she has acquired thus far; p has simply not felt the impact of her evidence. While this way (p.61) of understanding a subject's lack of evidence is incomplete in some respects (for instance, it so far counts a rational subject with a great deal of evidence that merely confirms, but does not change, her initial credence as having no evidence), it is a fine start. It can capture an important way of understanding a subject's lack of relevant evidence about whether the 2018 Olympics will be in South Korea, whether amber is heterogeneous in composition, whether someone who played on the Michigan football team in the late 1800s became chairman of the Iowa Democratic Party, or whether electric potential has only magnitude and not direction.

My claim is that when a subject is in these sorts of impoverished evidential circumstances with respect to an ordinary contingent p, she is epistemically permitted to suspend judgment about p (the reader can assume that the absence of evidence is as transparent as can be to the subject). The dictum to respect one's evidence is fundamental for traditional normative epistemology. The absence of evidence norm is among the most minimal ways for a subject to respect her evidence. It says just (and very roughly) that when she

has none, she may suspend. When your evidence gives you no guidance as to whether an ordinary contingent proposition p is true or false, the norm says that you are thereby permitted to suspend judgment about p. The alternative is to say that (at least in some cases) despite having no evidence relevant to whether that sort of proposition is true or false, suspension of judgment is epistemically prohibited. It is extremely difficult to see how that could be right. In fact it is hard to think of evidential circumstances more appropriate for suspension about these sorts of propositions than these kinds of absolutely impoverished ones. If you are going to have some attitude towards an ordinary contingent proposition that you understand, but about which you have absolutely no evidence either for or against, you cannot be going wrong by suspending judgment. These are exactly the sorts of circumstances suspension of judgment is for.[12]

(p.62) This is not to say though that these are the only sorts of evidential circumstances in which suspension of judgment is epistemically permitted. That is, it is not only in absolutely impoverished evidential circumstances that subjects are permitted to suspend. For instance, subjects who have some evidence, but (say) not nearly enough to fully settle the question of whether p is true or false, are also epistemically permitted to suspend judgment. As are subjects with massive bodies of evidence that are (roughly) evenly balanced (as just discussed). Sextus may have thought that suspension was always rationally permissible.[13] This is fertile territory. For now I want to try to stay focused on absolutely impoverished evidential circumstances since they make for the tidiest versions of some of the arguments to come.

Here is how I will proceed. In section 2 I will argue that thinking about cases in which subjects lack evidence shows that suspension about long conjunctions and disjunctions as well as the individual conjuncts/disjuncts is epistemically permitted. This will leave the Straightforward Reduction in a perilous position. In section 3 I will buttress and extend these conclusions by thinking about priors. I will show that the Straightforward Reduction really cannot succeed and that suspension of judgment about $p,\neg p$ cannot (just) be a matter of having standard credences for p, $\neg p$ (no matter which). In section 4 I will suggest a way to extend the arguments from the earlier sections. And in section 5 I will tie things up and look further forward.

2. Conjunctions and Disjunctions

Say S is given a drawing of a snowflake and is told that real snowflakes from different locations will be collected. She is asked to consider, for each collected snowflake, whether that snowflake has the same structure or shape as the one in her drawing. S will not see any of the collected snowflakes. Let a_1 be the proposition that the first snowflake is a match, a_2 the proposition that the second snowflake is a match, and so on through to a_n (for some finite n). This S does not know very much about snowflakes, in fact she's never seen any before (she's from a small island in the tropics). She knows roughly what they are—that they are bits of ice that fall from the sky when it is cold—but other than that, she has no evidence at all about their shapes: she has no idea what sorts of shapes they can take, how many different shapes they come in, the frequency with which a given shape occurs, or occurs with other shapes, and so on. There is simply nothing in her body of total evidence that bears on whether the collected flakes match her flake drawing.

(p.63) Take a_1 – a_3. Say S considers each in turn. Is it epistemically permissible that she suspend judgment about each of these propositions individually? These are ordinary contingent propositions and we are stipulating that S has no evidence either for or against any of them. From this perspective, as we have seen, suspension is epistemically permissible. When she wonders whether a given flake is a match and recognizes her complete ignorance, she is permitted to suspend.

What about the conjunction of these propositions? Or the disjunction? Can S rationally suspend about those alongside suspending about a_1 – a_3 individually? It looks that way. If we're to imagine that S has no evidence at all that is relevant to whether each of a_1 – a_3 is true or false, and that she has no other evidence relevant to whether the conjunction or disjunction of those propositions is true or false, then it looks as though she also has no evidence relevant to whether that conjunction and disjunction are true or false. Given this, suspension of judgment still looks epistemically permissible. From the perspective of her evidence she is as ignorant about the conjunction and disjunction as she is about the individual conjuncts/disjuncts. But this means that we want to be able to say that S can suspend about a_1 – a_3 and about the conjunction and/or disjunction of those propositions at the same time.[14]

It looks as though the reasoning in the last paragraph carries over to longer conjunctions and disjunctions as well. For instance, imagine that on top of a_1 – a_3 S also considers a_4 – a_{10}. The very same

considerations about the epistemic permissibility of suspending about each individual a_i ($1 \leq i \leq 10$) apply. Again, if we assume that S has no evidence relevant to each and no other evidence relevant to conjunctions and disjunctions of these propositions, then she still has no evidence relevant to whether those conjunctions and disjunctions are true or false. Again, suspension about each of $a_1 – a_{10}$, as well as conjunctions and disjunctions of those are all epistemically permitted at a time. And the same goes for much longer conjunctions/disjunctions even. She might consider whether the first 100 flakes are a match, or the first 1,000, and so on. Our reasoning about $a_1 – a_3$ carries over. With no evidence for or against any individual conjunct/disjunct, and no additional evidence that bears just on the conjunctions/disjunctions, she has no evidence for or against those conjunctions/disjunctions either. In this sense, she has no idea about whether all or any of those flakes are a match. So now it looks as though it can be epistemically permissible for S to suspend judgment about very long conjunctions and disjunctions as well as each of their conjuncts/disjuncts.[15]

(p.64) This means trouble for the Straightforward Reduction. The easiest way to see this is to start with a plausible credence-theoretic account of agnosticism (for the Straightforward Reductionist) and see why it fails. The Straightforward Reductionist is going to try to find some subinterval of [0,1] that is such that suspending judgment about p, $\neg p$ is just a matter of having standard credences for p, $\neg p$ in that subinterval. Let's call this special subinterval the SJ-*subinterval*. Having credences for p, $\neg p$ in the SJ-subinterval is necessary and sufficient for being in a state of suspended judgment about p, $\neg p$ for the Straightforward Reductionist. Presumably, the first-pass suggestion is that the SJ-subinterval should be a roughly middling subinterval of [0,1]. Earlier I said that the agnostic attitude is an attitude of neutrality or indecision about the truth of p, and a middling p-credence surely looks like a good candidate for capturing that sort of neutrality or indecision. Moreover, the Straightforward Reductionist is going to have to make room for belief and disbelief in the formalist's taxonomy and presumably those will take up the high and low subintervals of [0,1] (respectively). Let's start with the suggestion that the SJ-subinterval is [1/3,2/3] (nothing hangs on this specific precisification of 'middling'; in a moment we'll see why others fail as well).

(Mid$_1$) S suspends about p at t iff $C_s(p) \in [1/3,2/3]$ at t.

But the Straightforward Reductionist has to say that **(Mid $_1$)** is false. It is false because it renders impermissible combinations of suspendings that as we've just seen need not be impermissible. To see this we need only focus on $a_1 - a_3$. Given that a_1, a_2, a_3 are probabilistically independent, $C_s(a_1 \wedge a_2 \wedge a_3)$ ought to be equal to $C_s(a_1)C_s(a_2)C_s(a_3)$.[16] But if each of $C_s(a_1), C_s(a_2)$ (p.65) and $C_s(a_3)$ are middling$_1$ (in [1/3,2/3]), then $C_s(a_1)C_s(a_2)C_s(a_3)$ cannot be middling$_1$, and so $C_s(a_1 \wedge a_2 \wedge a_3)$ ought not to be middling$_1$.[17] And $C_s(a_1 \vee a_2 \vee a_3)$ ought to be equal to $C_s(a_1) + C_s(a_2) + C_s(a_3) - C_s(a_1 \wedge a_2) - C_s(a_1 \wedge a_3) - C_s(a_2 \wedge a_3) + C_s(a_1 \wedge a_2 \wedge a_3)$. But if each of $C_s(a_1), C_s(a_2)$ and $C_s(a_3)$ are middling$_1$, then $C_s(a_1) + C_s(a_2) + C_s(a_3) - C_s(a_1 \wedge a_2) - C_s(a_1 \wedge a_3) - C_s(a_2 \wedge a_3) + C_s(a_1 \wedge a_2 \wedge a_3)$ cannot be middling$_1$, and $C_s(a_1 \vee a_2 \vee a_3)$ ought not to be middling$_1$.[18] If S's credences for $a_1 - a_3$ are middling$_1$ at t, then she is not permitted to have a middling$_1$ credence for the conjunction or for the disjunction of those propositions at t. If **(Mid $_1$)** is true then S can only suspend about each of $a_1 - a_3$ as well as their conjunction or disjunction at t by having credences she is not permitted to have at t. **(Mid $_1$)** renders the relevant combinations of suspendings epistemically impermissible. But we just saw that these combinations of attitudes are permissible. **(Mid $_1$)** is false.

In general, if we assume that credences are standard, then if **(Mid $_1$)** is true it is never epistemically permissible for S to suspend judgment about each of (at least) three probabilistically independent propositions as well as the conjunction of those three propositions, and it is never epistemically permissible for S to suspend judgment about each of (at least) three probabilistically independent propositions as well as the disjunction of those three propositions. This does not look like a good result.[19] Either way, we have now confirmed that it is false: this combination of suspendings can indeed be epistemically permissible.

So, the Straightforward Reductionist cannot say that the SJ-subinterval is a middling$_1$ subinterval of [0,1]. But we also know that this is only the tip of the iceberg. Even if we focus on just $a_1 - a_{10}$ and say that $C_s(a_i) = 0.5$, then $C_s(\wedge_{i=1}^{10} a_i)$ ought to be 0.0009765625. And even this is still just the tip of the iceberg. Given the sort of reasoning from the snowflake case, it doesn't look as though there is any n such that suspending about ($\wedge_{i=1}^{n} a_i$) alongside each individual conjunct is epistemically impermissible (and the same goes, *mutatis mutandis* for the disjunctions). If this is right then take any x, y such that $0 < x \leq y < 1$, and say that the SJ-subintervals [x, y]. No matter what x and y are we can extend our snowflake example so that there

are conjunctions and disjunctions such that it is epistemically permissible for S to suspend about those conjunctions and disjunctions as well as their (p.66) individual conjuncts/disjuncts at t, but it is not epistemically permissible for her to have credences for all of these in [x,y] at t. The only proper subinterval of [0,1] that won't be susceptible to this sort of counter-example will be the open unit interval (0,1).[20] If S's credences for all of the a_is are in (0,1) then her credences for the conjunctions and disjunctions ought to be as well.[21] So far then it looks as though the Straightforward Reductionist has to say that the SJ-subinterval is at least (0,1).[22]

There might be some temptation to say that S shouldn't have any credence at all in these snowflake propositions. After all, she's completely ignorant about snowflake shapes. But we know that it is epistemically permissible to suspend about the relevant propositions. If it's permissible to suspend about p but not permissible to have a standard credence for p, then suspending about p cannot (just) be a matter of having a standard credence for p. In general, if it's possible for S to suspend about these snowflake propositions (which it is) but to not have credences for them, then any position that aims to reduce suspension to having a credence in some range (standard or otherwise) is in some trouble.

These sorts of "credence gaps" are a real worry for the Straightforward Reduction then. If the a_is don't generate credence gaps we can try to think of other propositions and circumstances that might. A subject might be so utterly in the dark about what the price of copper will be in 100 years, or whether in Minuscule 545 iota adscript occurs up to Luke 1:77, then ceases, or whether the Hill 50 Gold Mine was Australia's most profitable mine between 1955 and 1961 that he ought to simply refuse to have any degrees of belief for the relevant propositions. But it is epistemically permissible that he suspend judgment about the relevant propositions even if we think he can't or shouldn't assign credences to them. No matter how confused or ignorant one is about how much credence to give a proposition in a case, so long as one can grasp the proposition in that case, suspension looks epistemically appropriate.[23]

(p.67) So far it looks as though the Straightforward Reduction has to have the entire open unit interval dedicated to suspension of judgment, leaving 1 for belief and 0 for disbelief. I take it this is a bad result. We want to be able to say that subjects can (rationally) believe p with a p-credence less than 1. Certainly if credences are closely related to betting dispositions in the ways that they are standardly thought to be, that is, as identical to dispositions to bet or as

explanations of such dispositions, then there is next to nothing that we believe to degree 1. Even if we wanted to keep credences and bets apart, we should still be troubled by the association of belief with credence 1. At least the following is true of credence 1: it cannot be moved by conditionalization. A *p*-credence of 1 ought to remain 1 no matter what evidence comes in, given the standard way of updating a credence function on new evidence. It is effectively rational to hold on to a belief with credence 1 come what may. Perhaps there are some beliefs like this, but that class is extremely limited. Whatever notion of belief the Straightforward Reduction is left with here, it is not our everyday notion: there is next to nothing we believe by the lights of the Straightforward Reduction.[24] Those less worried about the identification of belief with credence 1 and disbelief with credence 0, can take only momentary comfort: in the next section I'll show why the Straightforward Reduction doesn't even have this option.

3. Priors

I am going to assume that the Straightforward Reductionist is a Bayesian, at least in the following additional sense: she thinks that the rational updating of a subject's credence for *p* upon receiving new evidence *e* ought to be determined in part as a function of her prior credence for *p*—her credence for *p* "before" she received *e*. Usually this prior credence is the result of updating on past evidence (i.e. it was the posterior credence in some other update), (p.68) but this need not be the case. The Bayesian makes room for an absolutely prior *p*-credence: the subject's credence for *p* that is prior to the receipt of any evidence for or against *p* at all. This sort of primordial or original prior credence is sometimes called an ur-prior credence. There is nothing particularly mysterious about ur-priors. In order to have our opinions rationally informed by evidence, we have to start with some prior, uninformed state of opinion. We can think of one's credence for *p* at any time *t* as the one that results from conditionalizing this ur-prior for *p* on one's total evidence at *t*. One's ur-prior *p*-credence is the credence for *p* one has in the complete absence of evidence: it is an a priori state of opinion about *p*; it is a degree of belief assigned to *p* a priori.

In this section I'll argue—by thinking about ur-priors—that the Straightforward Reductionist must make the SJ-subinterval [0,1]. First, I'll say why, if credences are standard, it can be epistemically rational to have ur-priors anywhere in [0,1]. After that I'll draw out the implication of that fact for the Straightforward Reduction.

Assume credences are standard and imagine a subject S assigning ur-priors to the propositions in various finite partitions $P_i (i = 1, 2, \ldots, n)$, where P_1 is a one-celled partition, P_2 a two-celled partition, and so on. These cells are propositions. What should a rational S's ur-prior credences over these partitions look like? We need only be concerned with the minimal demands that her credence for each cell/proposition be in [0,1] and that her credences over P_i – for any i – sum to 1.[25] These demands guarantee that however a rational S distributes her credences over the cells of the P_is, so long as there are sufficiently many of them, her credences for most of those propositions will have to be very low (and her credences for their negations very high).[26] As i increases, a rational S will have to have credences for the cells that approach 0 (and credences for their negations that approach 1).

Moreover, if we move to thinking about partitions with infinitely many cells, and keep in place the Straightforward Reductionist's assumption that credences are standard, we can see that a rational S will have to be able to (p.69) have ur-priors that are 0 or 1 as well. Say Q is such an infinite partition. If Q is merely countably infinite, then a rational subject can have a credence greater than 0 for each $q \in Q$, but her credences over Q cannot be uniform. When the number of possibilities is countably infinite, a rational subject with standard credences must favor some possibilities over others.[27] But when Q is uncountable, if credences are standard, then S will be rationally required to have credence 0 for uncountably many members of Q (and credence 1 for each of their negations). If S's credences over an uncountable Q are to sum to 1, then she can only have credence greater than 0 for countably many $q \in Q$. But this means that a rational S will have credence 0 for uncountably many $q \in Q$ (and so credence 1 for their negations).[28]

What impact should this have for the Straightforward Reduction? In section 2 I argued that the Straightforward Reductionist has to say that the SJ-subinterval is at least (0,1). In fact, let's imagine a Straightforward Reductionist who says just that: suspending about p, $\neg p$ is just a matter of having credences for p, $\neg p$ in (0,1), believing p just a matter of having credence 1 for p, and disbelieving p just a matter of having credence 0 for p. I have already said a little bit about why this is an unpalatable view, but now we have a new line on how it goes wrong. Think about our uncountable partition Q, and rational S distributing her credences over Q a priori. Let's say Q is a partition of ordinary contingent propositions. The Straightforward Reductionist I just described will have to say that a rational S disbelieves

uncountably many $q \in Q$. He has to say that S can suspend about every $q \in Q$ only by having probabilistically incoherent credences. So he has to say that suspending about every $q \in Q$ is epistemically impermissible. But each $q \in Q$ is an ordinary **(p.70)** contingent proposition about which S has absolutely no evidence, and given that, suspension about every $q \in Q$ should be epistemically permissible. This means that the suggestion that the SJ-subinterval is (0,1) renders rationally impermissible combinations of suspendings that are not rationally impermissible. To avoid this result it looks as though the Straightforward Reductionist will have to say that the SJ-subinterval is [0,1]. Let me flesh out and buttress this argument.

First, let me say a bit more in defence of the claim that suspension about these ordinary contingent partition propositions is permissible. Remember, in these cases, S's credences are had in the compete absence of evidence; her credences are ur-prior credences. This means that we are to think of them as opinions she has completely a priori. But now we can buttress the claim that suspension about these propositions is epistemically permissible as follows: it is not epistemically permissible that S believe or disbelieve an ordinary contingent proposition a priori, and so with respect to these partitions of ordinary contingent propositions, suspension is not only a rationally permissible doxastic attitude to have towards those partition propositions, it is the uniquely epistemically permissible attitude (from the traditionalist's taxonomy) to have. Let's flesh this out further.

Let's first focus on a single relevant (countable) partition A. We can think of A as an ordinary contingent question, in this case the question: how many birds are there in France? This question can be thought of as a partition with each possible complete answer (that there are no birds in France, that there is exactly one bird in France, that there are exactly two birds in France, and so on) as a cell of that partition. These answers are mutually exclusive and exhaustive. Now we can imagine S distributing her credences over A a priori. For any $a \in A$, it is not epistemically permissible that S believe a. I take it that this is obvious. It is not epistemically permissible for S to believe that there is some specific number of birds in France a priori. But what about disbelief? Is it also epistemically impermissible that she believe any $\neg a$ a priori?

On the one hand there is an obvious and important symmetry between belief and disbelief in this case. Any given $a \in A$ is an ordinary contingent proposition and so is its negation. It is standard to claim that it is not epistemically permissible to believe those sorts of propositions a priori. On the other hand, if credences are standard, then S's credences for most of the propositions in A ought to be extremely low, and from that perspective it can look as though there is also an important asymmetry between the epistemic permissibility of believing some a and the epistemic permissibility of believing some $\neg a$. If it is epistemically permissible that S be extremely confident that some $\neg a$ is true a priori, should it really be impermissible that she believe it a priori? While I admit that there is something compelling about this line of thought, there is near consensus that ordinary contingent propositions should not be believed a priori. There have been a few arguments over the years that some contingent truths are knowable a priori. Kripke (1980) famously argued that someone who introduces the term 'one meter' as a rigid designator for the (p.71) length of a particular stick s at time t can know a priori that the length of s at t is one meter. These "superficially contingent" propositions are such that the subject has some actual semantic guarantee of their truth, and so it looks as though she is permitted to believe them a priori.[29] Some have even argued that there may well be some deeply contingent propositions that we can be permitted to believe a priori.[30] But these too form a limited class (e.g. there is a believer), and the as do not fall into it. This is all to say that there is almost no precedent for claiming that propositions like the $\neg a$ s could be rationally believed a priori.[31] Insisting otherwise here would at the very least amount to a significant re-drawing of the bounds of the a priori, and without serious argument for that re-drawing, we should disregard the relevant asymmetry between any a and $\neg a$ here. In general then we should conclude that neither belief nor disbelief towards any $a \in A$ is epistemically permissible.

A is a countable question or partition, but everything I have said about it will carry over to uncountable ones as well. S might contemplate questions with uncountably many answers: what the president's credence that it will rain tomorrow is, or exactly how long the tail of the oldest cat in the world is, or what the exact landing point of the winning dart at the 1983 BDO World Championships was, and so on. Everything I have said about A applies to these partitions as well. No matter how many cells make up these partitions, so long as they are partitions of ordinary contingent propositions neither belief nor disbelief is epistemically permitted.

This leaves suspension as the uniquely epistemically permissible attitude from the traditionalist's taxonomy for S to have towards these propositions. I don't think that the permissibility of suspension in these sorts of cases is up for grabs. I have argued that belief and disbelief are not epistemically permissible attitudes to have towards these propositions given S's evidential circumstances. S is epistemically prohibited from believing ordinary contingent propositions a priori. But these considerations do not extend to her suspending about these propositions. In fact, suspension looks like exactly the right attitude to have if she is going to have one (and there is no special reason to demand that she have none). If she's considering a priori whether apples are sweeter than spinach or how long it takes beavers to build dams or (p.72) where B. B King was born, then it is perfectly reasonable for her to suspend judgment.[32]

Now we can be clear about exactly what this all means for the Straightforward Reduction. The Straightforward Reductionist wants to isolate an SJ-subinterval—a subinterval of [0,1] such that suspending about p, $\neg p$ is just a matter of having credences for p, $\neg p$ in that subinterval. We have already seen why anything short of (0,1) won't work: picking any subinterval shorter than that will force the Straightforward Reductionist to say that some epistemically permissible combinations of suspendings are not epistemically permissible (e.g. long conjunctions and each of their conjuncts). This section confirms that result, but also extends it.

It confirms it as follows. Say the Straighforward Reductionist sets the SJ-subinterval to some interval [x,y] (or (x,y)) shorter than (0,1). We can now find a finite partition of ordinary contingent propositions such that S is permitted to suspend judgment about all of the propositions in that partition, but such that her credences for at least some of those propositions ought to be less than x and her credences for their negations greater than y. But this means that the cases in this section again show that if the Straightforward Reductionist sets the SJ-subinterval to anything shorter than (0,1) she will render epistemically impermissible combinations of suspendings that are not epistemically impermissible. But I hope it is also clear now that (0,1) suffers in much the same way. If the SJ-subinterval is (0,1), then it is not rationally permissible that S suspend about the ordinary contingent propositions in some infinite partitions. But suspending about all of those propositions is epistemically permissible. Moreover, it looks as though the Straightforward Reductionist who makes the SJ-subinterval (0,1) (or shorter) will have to say that a rational S believes many of those ordinary contingent propositions a priori. But

a rational *S* won't believe ordinary contingent propositions a priori. If the SJ-subinterval is (0,1) (or shorter), then what look like perfectly rational ur-prior credence distributions over these infinite partitions are rendered irrational since they will have the result that the relevant subject believes ordinary contingent propositions a priori.

So the SJ-subinterval cannot be (0,1) either. In fact, it cannot be any proper subinterval of [0,1]. The only option left for the Straightforward Reductionist is to make the SJ-subinterval [0,1] itself. The Straightforward Reduction has the result that the SJ-subinterval is [0,1]. But this is not an acceptable result for the Straightforward Reduction. It says that suspending about *p* is just a matter of having any degree of belief for *p* at all. But if suspending about *p* is a matter of having any standard *p*-credence at all, then there is no hope of reducing believing *p* and disbelieving *p* to having standard *p*-credences: the whole unit (p.73) interval has to be devoted to suspension of judgment. The Straightforward Reduction cannot succeed since once suspending about *p* is just a matter of having a standard credence for *p*, believing *p* cannot be a matter of having a standard credence for *p*, and disbelieving *p* cannot be a matter of having a standard credence for *p*.

Moreover, the conclusion that *S* suspends about *p*, ¬*p* iff *S* has (any) standard credences for *p*, ¬*p* is not a palatable conclusion even for someone not concerned with reducing believing *p* and disbelieving *p* to having standard credences for *p* (and ¬*p*). We are not agnostic about *p* simply in virtue of having a credence for *p*. This would not only make it that believing *p* and disbelieving *p* are not reducible to standard credences for *p* but that it is irrational to either believe *p* or disbelieve *p* when one has a standard credence for *p*. If anyone with a standard credence for *p* is agnostic about *p*, then on the assumption that believing *p* and suspending about *p* or disbelieving *p* and suspending about *p* are irrational combinations of attitudes, no one with a standard credence for *p* is permitted to believe *p* or disbelieve *p*; that is, no one is permitted to be in a state in which they both have standard credences for *p*, ¬*p* and either believe *p* or disbelieve *p*.[33] We should conclude not only that the Straightforward Reduction fails, but also that suspension about *p* is not (just) a matter of having a standard credence for *p*.

4. Beyond the Straightforward Reduction

The Straightforward Reduction fails. Believing *p*, disbelieving *p*, and suspending about *p* cannot all just be matters of having standard credences for *p*. The endeavour falls apart once we try to reduce suspension of judgment about *p* to a standard credence for *p*. More generally, we should conclude that suspending judgment about *p* is not (just) a matter of having a standard credence for *p*. Do my arguments here leave room for a position like Standard Lockeanism though? The Lockean is someone who thinks that believing *p* is just a matter of having a sufficiently high degree of belief for *p*, and the Standard Lockean is someone who thinks this about belief and wants to maintain that credences are standard. The Standard Lockean won't be bothered if being in a state of suspended judgment about *p*, ¬*p* cannot be reduced to having standard credences for *p*, ¬*p*, but he does want to maintain that credences are standard and that (dis)believing *p* is reducible to having a standard credence (p.74) for *p*. I want to make clear that my arguments here do not leave room for this Standard Lockean.

In this paper I have described a number of cases in which suspending judgment about *p* is epistemically permissible. Giving up on reducing suspending about *p* to a standard credence does nothing to change this. The Standard Lockean (and anyone else) will still need to make room for rational suspension of judgment in those cases. But this will not be possible if the Standard Lockean insists that rational subjects in the sorts of cases I have been describing have standard credences. If they do still have standard credences then they will have very high/ low credences. Given this, the Standard Lockean will have to say that they are believers and disbelievers, and so cannot be rationally agnostic. This Lockean has to say that these subjects are not permitted to suspend in the complete absence of evidence and that they are permitted to believe ordinary contingent propositions a priori.

Put slightly differently, giving up on reducing suspension about *p* to having a standard credence for *p* but maintaining that subjects in the relevant cases have standard credences, while respecting the norms for suspension and a priori believing I've discussed here, will have the result that suspension about *p* is rationally compatible with any standard credence for *p* (even if it is not reducible to having a standard credence). But if suspending about *p* can be rationally compatible with having any standard credence then no standard credence can amount to belief since believing *p* and suspending about *p* are not rationally compatible. So the Standard Lockean must do more than just give up on reducing suspending about *p* to having

some relevant standard credence for *p*. In fact, his only option is to claim that subjects in the cases I've described have no credences at all. While this might look plausible in some of those cases, it cannot be right for all of them. That would amount to claiming that rational subjects largely have no absolutely prior credences. If subjects need priors to update their degrees of belief on new evidence, then insisting that rational subjects have no absolutely prior degrees of belief, threatens to leave them with no posterior credences either, that is, threatens to leave rational subjects with no credences for anything at all. Obviously this is not a real option for the Standard Lockean, and so we must conclude that Standard Lockeanism is false. Anyone keen on reducing belief and disbelief to high and low credences (respectively), who wants to say that subjects in my cases do have degrees of belief, will have to allow for non-standard degrees of belief in those cases.

In fact, some (not just those keen on reducing belief to credence) argue that standard credences are not epistemically appropriate in some of the sorts of cases I have described. For instance, some claim that we need to use infinitesimal degrees of belief to accurately capture a subject's state of opinion when those opinions range over uncountably many possibilities.[34] (p.75) Making a credence function a hyperreal-valued function might help to escape the conclusion that suspending about *p* is rationally compatible with having a *p*-credence of 0 or 1, but the allowance will do nothing to avoid the result that suspension is rationally compatible with all real numbers (and now perhaps infinitesimals as well) between 0 and 1. Another non-standard approach might be more helpful in this regard though.

Some claim that in the absence of evidence, or in a state of complete ignorance one ought to have vague or imprecise (mushy, interval) credences. These are degrees of belief that are measured with or represented by sets of real numbers (e.g. [0.2,0.8]) rather than single real numbers.[35],[36] Can the Lockean (not the Standard Lockean) say that subjects in the cases I've described have imprecise credences?

Perhaps, however, this solution does not come all that easily.[37] Even in the cases I've described, the claim that rational subjects cannot have standard credences is a significant one. Someone who wanted to avoid the conclusion that suspension is rationally compatible with most any precise degree of belief by turning to imprecise credences will have to then say that a rational subject mostly does not have absolutely prior precise credences that are bound by the standard axioms of the probability calculus. This isn't just the claim that some rational subjects sometimes have or can have non-standard

credences. It's the claim that it is largely rationally impermissible to have standard ur-priors (at least for ordinary contingent propositions). While many Bayesians are comfortable with imprecise credences it isn't clear that everyone will be happy with this strong claim about priors (certainly the ultra-subjective Bayesian won't be, but it is easy to imagine complaints from the more objectively minded as well).

The much bigger problem though is that this "solution" is not a general solution at all. It may help with the cases I have discussed so far, but that is (p.76) not enough to guarantee that suspension of judgment about p is not rationally compatible with most any standard credence for p. That is the conclusion that the Lockean (and others like him) is trying to avoid. In particular, his goal is to have it that suspension about p is not compatible with high and low standard credences for p. While arguing that S ought not have standard credences in the complete absence of evidence might help to block one path to the conclusion that suspension about p is compatible with high and low p-credences, it does nothing to block other very nearby paths (nor further away ones, of course). Let's stay focused on questions about the epistemic permissibility of suspending with high or low credence. As I mentioned at the outset, the absence of evidence norm is just one norm among many for epistemically rational suspension. Even if we say that it needn't license suspension about p with high/low credence for p, that does not mean that some other (closely related) norm won't. It is not at all clear that suspension about p is going to be epistemically prohibited in cases in which it is relatively uncontroversial that p-credences ought to be precise and either very high or very low. Genuinely making the case that such a combination of attitudes is epistemically permissible is beyond the scope of this paper. But I want to end this section by saying a little bit in defence of the claim.

When chances are known it is fairly widely agreed that subjects should set their credences to match those known chances.[38] Let's imagine a slight variation on our French birds case. Say France is birdless at t_1, but at t_3 God will give the French some birds. At t_1 S assigns her ur-priors to propositions of the form 'the number of birds in France at t_3 will be x' (where $x \in N$). At t_1 she has absolutely no evidence that bears on how many French birds there will be at t_3. Given the arguments so far we should say that suspension about these French bird propositions is epistemically permissible at t_1. But now say that at t_2 God tells S that at t_3 there will be between 1 and 10^8 birds in France, and that the number of French birds will be

decided by a random process such that each number in the relevant range has the same probability as any other (maybe God will roll his 100,000,000-sided die). Say B = $\{b_1, b_2, \ldots, b_{100,000,000}\}$ is the partition of French bird number propositions (where b_1 is the proposition that there will be exactly one bird in France at t_3, and so on). It is fairly widely agreed that at t_2 S ought to have the same, precise degree of belief for each $b_i \in B$, that is, for every $b \in B$, $C_s(b)$ ought to be $1/10^8$.

The challenge now is to say just why this sort of information should make suspension of judgment about the number of birds there will be in France at t_3 epistemically impermissible. *S* has more information than before about the number of French birds at t_3, and so we may not be able to lean as much on worries about the contingent a priori, but the new information that she (p.77) has is still incredibly minimal. While she knows how many possibilities there are, and that each is equally probable, she knows nothing at all that makes any remaining possible answer to the question of how many birds there will be in France at t_3 stand out from any other. That is, with respect to the question of the number of birds there will be in France at t_3, the evidence she has does not discriminate between the remaining possible answers. Why should suspension of judgment be epistemically impermissible when one's evidence does nothing to support any one possible remaining outcome over any other?[39]

We can make the worry acute as follows. Think about the change in view that *S* undergoes when she learns that the possible number of birds in France at t_3 is limited to within a certain range and will be determined by a chance process. Let's call this new evidence *e*. *e* is evidence in favor of the propositions in B. It is evidence that one of those propositions is true, and that the French bird propositions not in that set are false. Here is one way of fleshing this thought out. Before acquiring *e* let's say that *S* had credences for all of the French bird propositions, (for each $n \in N$ there is one such proposition). We can assume that *S* is rational and updates her credences as she ought. Given that, when *S* learns *e* her credence for any b_k (that the number of birds in France at t_3 will be k) for every k greater than 10^8 will drop to 0 (as will her credence that there will be no birds in France at t_3). What should we say about her confidence in the remaining 10^8 French bird propositions? Well, surely her confidence in at least some of them will increase upon learning *e* and dropping her credence in those French bird propositions incompatible with *e*. Even if one wants to insist that S's confidence is not to be measured with standard credences before acquiring *e*, we want to be able to say that rational *S* is more confident about at least some of the remaining propositions

after discovering *e* than she was before discovering *e*. Her outcome space is significantly smaller and more stable and so her confidence in at least some of the remaining propositions should have increased. However we cash out this increase in confidence (it won't be straightforward if we're comparing a non-standard credence to a standard one), the point is that we want to be able to somehow say that *e* makes it epistemically permissible that *S* be more confident about at least some of the remaining answers than she was before receiving that evidence. *e* is evidence in favor of those propositions.

Let's say that b_5 is one of the propositions about which *S* becomes more confident upon receiving *e*. We know that upon receiving *e* (at t_2), *S* credence for each $b \in B$ ought to be $1/10^8$. Anyone keen on avoiding the conclusion (p.78) that suspension of judgment is rationally compatible with high/low credences will have to say that suspension about b_5 (or any other $b \in B$) is no longer permissible at t_2. But is that plausible? Obviously believing b_5 (or any other $b \in B$) is not epistemically permissible at t_2. So what can we say about *S*'s doxastic options if suspension is impermissible (and so is belief)? There are two possibilities: either *e* permits disbelief (belief in $\neg b_5$) or it does not. If we say that disbelief is also epistemically prohibited, then we are left having to say that in acquiring good new evidence, evidence that grounds greater confidence in b_5, and a clear and precise credence for that proposition, *S* is no longer permitted to have any attitude from the traditionalist's taxonomy towards b_5. This can't be right. Moreover, this option obviously will not sit well with the Lockean who wants to identify disbelief with low credence. But if we say that disbelief is permitted, then we'll have to say that in acquiring good new evidence for b_5, evidence that grounds greater confidence in b_5 (but not belief), the uniquely permissible attitude for *S* to have towards b_5 becomes disbelief. We will have to say that in acquiring evidence that counts in favor of b_5 and grounds increased confidence in that proposition (but not belief), *S* will move from circumstances in which suspension is permitted to ones in which only disbelief is. But this can't be right either. Say I am suspending about whether my national team will win some Olympic event. I have no evidence either way. Then I learn that a member of one of the the opposing teams has a minor injury. This is evidence that counts in favor of my team winning; it grounds increased confidence that my team will win, but cannot ground a belief that they will win (the injury is minor, there are still plenty of other teams in the competition). It would be absurd to claim that I am now no longer permitted to suspend judgment about whether my team will win, but only permitted to believe that my team won't win: I've just acquired more evidence that they will

win; I am rationally more confident than before that they will win. But the same is true when it comes to S, e, and b_5. Given that at t_2 S has more evidence that there will be exactly five birds in France at t_3 than she did at t_1, it would be bizarre to say that at t_2 her evidence demands she move from suspension to disbelief about that proposition. It simply does not look as though this new evidence can bar suspension of judgment; suspension should still be epistemically permissible. If this is right then we get the result that suspension about p is rationally compatible with very high and very low credence for p.

This is just the start of a new argument for the rational compatibility of high/low standard credences for p with suspension about p. It extends the sort of general worries that have emerged in this paper about minimal evidence justifying maximal credence alongside suspension of judgment, but it requires further thinking about the conditions under which suspension of judgment is epistemically permissible (among other things).[40] And of course making the (p.79) case that suspension about p is rationally compatible with high/low credences genuinely compelling would involve much more discussion of the roles of credences, belief, and suspension of judgment in cognition, inquiry, and action more generally.

5. Concluding Remarks

Where does this leave us? I have argued that the Straightforward Reduction fails. Given that suspending judgment about ordinary contingent propositions is epistemically permissible in the absence of evidence, the demand that suspending about *p* be a matter of having a standard credence for *p* has the result that suspending about *p* is a matter of having any standard credence for *p*. If that's right then believing *p* cannot be a matter of having a standard *p*-credence and disbelieving *p* cannot be a matter of having a standard *p*-credence. However, the aim of the Straightforward Reduction is to reduce all three attitudes from the traditionalist's doxastic taxonomy to some relevant degrees of belief. But it looks as though the only way to reduce suspension about *p*, ¬*p* to some standard credences for *p*, ¬*p* is to make it that believing *p* and disbelieving *p* can no longer be so reduced. So the Straightforward Reduction cannot succeed.

Furthermore, the result that suspending about *p* is a matter of having any standard credence for *p* at all is not acceptable even for someone not concerned with reducing believing *p* and disbelieving *p* to standard credences for *p*. It looks false that anyone with a standard credence for *p*—any at all—is agnostic about *p*. Moreover, it would mean not only that (dis)believing *p* could not just be a matter of having a standard credence for *p*, but that (dis)believing *p* couldn't even be rationally compatible with having a standard credence for *p*. We get this last result since both (dis)believing *p* and suspending about *p* is not a rational combination of attitudes at a time, and so if anyone with any credences at all for *p*, ¬*p* at *t* is suspending about *p*, ¬*p* at *t*, then no one with credences for *p*, ¬*p* at *t* can also rationally believe *p* or ¬*p* at *t*. We must conclude that suspending about *p*, ¬*p* is not (just) a matter of having standard credences for *p*, ¬*p*.

I have also argued that anyone who wants to try to avoid the result that suspension of judgment about *p* is (rationally) compatible with high/low credence for *p* will have to say that subjects in the cases I discussed in sections 2 and 3 have non-standard credences in those cases. I ended though by saying a little bit about why—even if subjects in those cases do have non-standard credences—the conclusion that suspension of judgment about *p* is genuinely (rationally) compatible with very high and very low credences for *p* might prove difficult to escape. If it does prove unavoidable, then common versions of Lockeanism are false. If we can rationally suspend about *p* despite having (say) a very high degree of belief for

p, then believing *p* must be more than or different from merely having a high degree of belief for *p*.

(p.80) What about our two taxonomies? Even if the Straightforward Reduction fails, there are plenty of more sophisticated ways we might try to reduce the traditionalist's doxastic attitudes to degrees of belief. For instance, perhaps the difference between believing *p* and suspending about *p* does not lie in one's unconditional *p*-credences but elsewhere in one's credence distribution: in some of one's conditional *p*-credences or even one's credences for propositions other than *p* (e.g. some have tied suspending about *p* to having higher-order beliefs about one's first-order epistemic standing). We might also try starting with a different, non-Kolmogorov axiomatization. And I think that reductions in the other direction (i.e. reducing degrees of belief to traditional doxastic attitudes) are more promising than is often acknowledged. Each of these suggestions has some plausibility but faces difficulties, and it remains to be seen whether any can succeed. For now we should conclude that the Straightforward Reduction fails, and that suspending about *p* is not (just) a matter of having a standard credence for *p*. And we should be worried about whether suspending about *p* isn't genuinely rationally compatible with most any precise degree of belief for *p*.

References

Bibliography references:

Avnur, Y. (2011). Hawthorne on the Deeply Contingent A Priori. *Philosophy and Phenomenological Research 83*(1), 174–83.

Bergmann, M. (2005). Defeaters and Higher-Level Requirements. *The Philosophical Quarterly 55*, 419–36.

Burge, T. (1993). Content Preservation. *Philosophical Review 102*(4), 457–88.

Burnyeat, M. and M. Frede (1997). *The Original Sceptics: A Controversy*. Indianapolis: Hackett Publishing Co.

Christensen, D. and H. Kornblith (1997). Testimony, Memory and the Limits of the a Priori. *Philosophical Studies 86*(1), 1–20.

de Finetti, B. (1970). *Theory of Probability*, Volume 1. New York: Wiley.

Douven, B. I. (2008). The Lottery Paradox and Our Epistemic Goal. *Pacific Philosophical Quarterly 89*(2), 204–25.

Elga, A. (2010). Subjective Probabilities Should Be Sharp. *Philosophers' Imprint 10*(05), 19, 1–11.

Empiricus, S. (1933). *Outlines of Pyrrhonism*. Loeb Classical Library. Cambridge, MA: Harvard University Press.

Fantl, J. and M. McGrath (2009). *Knowledge in an Uncertain World*. New York: Oxford University Press.

Foley, R. (1992). The Epistemology of Belief and the Epistemology of Degrees of Belief. *American Philosophical Quarterly 29*, 111–24.

Friedman, J. (2013). Suspended Judgment. *Philosophical Studies 162*(2), 165–81.

Ganson, D. (2008). Evidentialism and Pragmatic Constraints on Outright Belief. *Philosophical Studies 139*(3), 441–58.

Hájek, A. (1998). Agnosticism meets Bayesianism. *Analysis 58*, 199–206.

(p.81) —— (2003). What Conditional Probability Could Not Be. *Synthese 137*(3), 273–323.

Hawthorne, J. (2002). Deeply Contingent A Priori Knowledge. *Philosophy and Phenomenological Research 65*(2), 247–69.

Hawthorne, J. and L. Bovens (1999). The Preface, the Lottery, and the Logic of Belief. *Mind 108*(430), 241–64.

Joyce, J. M. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives 19*, 153–78.

Keynes, J. M. (1921). *A Treatise on Probability*. London: Macmillan.

Kripke, S. A. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In *Philosophical Papers: Volume II*. Oxford: Oxford University Press.

Monton, B. (1998). Bayesian Agnosticism and Constructive Empiricism. *Analysis 58*, 207–12.

Salmon, N. (1995). Being of Two Minds: Belief with Doubt. *Noûs 29*, 1–20.

Skyrms, B. (1980). *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. New Haven: Yale University Press.

Sturgeon, S. (2008). Reason and the Grain of Belief. *Noûs 42*, 139–65.

—— (2010). Confidence and Coarse-Grained Attitudes. In *Oxford Studies in Epistemology Volume 3*. Oxford: Oxford University Press.

Turri, J. (2011). Contingent A Priori Knowledge. *Philosophy and Phenomenological Research 83*(2), 327–44.

van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.

van Fraassen, B. C. (1998). The Agnostic Subtly Probabilified. *Analysis 58*, 212–20.

Weatherson, B. (2005). Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives 19*(1), 417–43.

White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives 19*(1), 445–59.

—— (2010). Evidential Symmetry and Mushy Credence. In *Oxford Studies in Epistemology Volume 3*. Oxford: Oxford University Press.

Williamson, J. (1999). Countable Additivity and Subjective Probability. *British Journal for the Philosophy of Science 50*(3), 401–16.

Williamson, T. (2007). How Probable is an Infinite Sequence of Heads? *Analysis 67*(3), 173–80.

Notes:

versions of many of the ideas here. Finally, I would like to thank the Social Sciences and Humanities Research Council of Canada for their generous support.

(^{1}) 'Suspending judgment about p' is ungrammatical if 'p' is to be replaced with a declarative complement. 'Suspend judgment about' can embed interrogative complements and noun phrases, but not declaratives. In these sorts of constructions (and some others throughout) 'p' should be replaced with sentences like, 'the proposition p'.

(^{2}) This is a very informal account of the relevant set-up, but it should do just fine. One thing worth flagging is that there are constraints placed upon the set of propositions that is the domain of a subject's credence function. Let's call this set F. F should contain all of the truth-functional combinations of its members: if p and q are in F, then $\neg p$, $\neg q$, $(p \wedge q)$, $(p \vee q)$, $(p \rightarrow q)$, and $(p \leftrightarrow q)$, and so on are all also in F. I will just assume that F has these properties, but one can think of them as normative constraints instead if one prefers.

(^{3}) These standard norms are:

> (3) *Non-Negativity*: $C_s(p)$ ought to be a non-negative real number.
> (4) *Normalization*: $C_s(p)$ ought to be 1 if p is a tautology ($\vDash p$).
> (5) *Countable Additivity*: for any countable sequence of mutually exclusive outcomes $p_1 \dots p_n \dots$, $C_s(\vee_j p_j)$ ought to equal $\sum_j C_s(p_j)$.

The basic picture is then one according to which a single real-valued function that is normatively bound by (3)–(5) represents a subject's total doxastic standing at a time.

(^{4}) Although, I will assume that suspension is closed under negation: that S suspends about p at t iff S suspends about $\neg p$ at t, and (for ease of exposition) that $C_s(p) = x$ iff $C_s(\neg p) = 1 - x$ (even if I sometimes do not say so).

(^{5}) For a small, but good discussion of suspension and degrees of belief, see van Fraassen (1998), Hájek (1998), and Monton (1998).

(^{6}) The Straightforward Reduction (or something quite like it) was previously besieged by the Lottery Paradox (see Foley (1992) for a good discussion). The worrying attack though relied on a closure principle for rational belief that some chose to give up or deny. If

anything, abandoning that principle has become more, rather than less acceptable with time. Some of the arguments here then can be thought of as a new (and more dangerous) attack on the Straightforward Reduction that requires no such principle.

(^{7}) For some agreement see (e.g.) Bergmann (2005) and Sturgeon (2010).

(^{8}) For more on these and other arguments for thinking that suspension of judgment is a proper doxastic attitude, see Friedman (2013).

(^{9}) Unless I say otherwise, when I talk about S's credence for p, I mean her unconditional credence for p.

(^{10}) For a good discussion of the latter, see Hawthorne (2002).

(^{11}) For instance, Keynes (1921) proposed that e is *irrelevant* to p (given background knowledge K) if the probability of p on K is the same as the probability of p on $K + e$, and relevant otherwise.

(^{12}) Does this mean that belief is not permitted in these sorts of evidential circumstances? In general, I do not think that claims about the epistemic permissibility of suspension of judgment should be thought to entail or imply or be easily translated into claims about the epistemic impermissibility of belief. The possibility of a sort of "weak permissivism" should be open: there may well be evidential circumstances that permit either belief or suspension of judgment. That is, given some evidential circumstances with respect to p, it may be epistemically permissible to either (say) believe p or suspend about p (see White (2005) for some related discussion and potential pitfalls though). But is belief epistemically permissible in the sort of cases I am interested in here? In some cases I will argue that it is not, but otherwise I want to remain as neutral as I can on the issue. Can a subject both believe and suspend at once? I think that once we think of belief and suspension as independent attitudes we should think that the combination of attitudes is possible at a time. I take it though that we'll want to say that a subject who has both attitudes at once is at least sometimes in a rationally conflicted state (much like the subject who believes both p and $\neg p$ at a single time). Of course much more would need to be said about the nature of the objects of the agnostic attitude to get clear about this rational incompatibility. See Salmon (1995) for a good discussion of cases in which there may be no rational failing in both believing p (or $\neg p$) and suspending about p

at a time. I am going to try to skirt these sorts of issues in the paper. The state of both believing p (or $\neg p$) and suspending about p, $\neg p$ at a time will be treated as an irrational state here.

(^{13}) Sextus claimed that we ought to be aiming for a kind of doxastic tranquility which can be arrived at by suspending judgment on all matters. See Empiricus (1933), Book I for his famous discussion of these issues. Of course, exactly what Sextus thought we did or ought to believe (rather than suspend about) is a matter of some controversy. See Burnyeat and Frede (1997) for some pillars of the debate.

(^{14}) I have argued that suspending about each conjunct/disjunct is epistemically permitted and that so is suspending about the conjunction and disjunction, but is it epistemically permissible to suspend about all of those at once? I don't see why not. S has no evidence either for or against any of these, their combinations are neither contradictory nor tautologous, and so at this point we have no reason to enact a ban on this combination of suspendings. In general, this looks like a perfectly epistemically appropriate (and commendable) combination of attitudes to have.

(^{15}) There's nothing special about this snowflake example either. We could replace the a_is with other propositions. For example, p_1: the peace lily is a member of the Araceae family, p_2: the magnolia is a member of the Araceae family, p_3: duckweed is a member of the Araceae family. Or, q_1: *that* star is a Population I star; q_2: *that* star is a Population I star; q_3: *that* star is a Population I star. Or, r_1: A has the DMD gene; r_2: B has the DMD gene; r_3: C has the DMD gene. Or, s_1: A went to the party; s_2: B went to the party; s_3: C went to the party. And so on. In each case we're to imagine that the relevant subjects don't have any relevant evidence.

(^{16}) Worry. If S ought to treat $a_1 - a_3$ as probabilistically independent doesn't that mean that she needs to have some information about their dependence relations and doesn't that amount to her having some evidence? I'm not convinced of either. That is, it may well be true that she ought to treat them as independent even if she has no information about dependence relations (she has to assume something about dependence relations to have all of the relevant credences), and it isn't clear that her having this information amounts to her having evidence in the sense at issue here. It is worth pointing out though that even if she did have that information and we did count it as evidence, this doesn't significantly change the impact of the case. Suspension about the individual a_is as well as their

conjunctions and disjunctions still looks epistemically permissible. More importantly though, I want to make clear that a version of this argument goes through no matter what we say that S ought or is permitted to assume about dependence relations. S is not only permitted to suspend about the conjunction of $a_1 - a_3$ but any conjunction that for each a_i has either it or its negation as a conjunct, e.g. that flakes 1 and 2 are matches, but not flake 3, that flake 1 is not a match, but flakes 2 and 3 are, and so on. Each such conjunction is a possible outcome of this flake-matching experiment, and S is epistemically permitted to suspend about which of those outcomes obtain for the very same reasons she is permitted to suspend about whether the outcome in which all flakes match obtains. But if there are n a_is then there are 2^n such outcomes/conjunctions (when $n < 2$ these outcomes are obviously not equivalent to conjunctions of distinct atomic propositions). So long as $n > 1$, whatever we say that S ought to or is permitted to assume about dependence relations, she will not be permitted to have middling$_1$ credences for all of these conjunctions. In general, for any n a_is, probabilistic coherence alone requires S to have credence $1/2^n$ or less for at least one conjunction that for each a_i has either a_i or $\neg a_i$ as a conjunct. But if $n > 1$, $1/2^n \notin [1/3, 2/3]$. More on this to come in the next section.

(^{17}) If $x \leq 2/3$, then $x^3 \leq 8/27$. And $8/27 < 1/3$.

(^{18}) If $x,y,z \geq 1/3$, then $[x + y + z - xy - xz - yz - xyz] \geq 19/27$. And $19/27 > 2/3$.

(^{19}) In fact, the situation is worse than this even. If credences are standard, then it is never epistemically permissible that S be agnostic about *just two* probabilistically independent propositions as well as a conjunction *and* a disjunction of those propositions (I leave the details to the reader). This cannot be right.

(^{20}) Of course, there are other options left, i.e. [0,1), and (0,1]. I am simply assuming that the best case scenario for the Straightforward Reductionist is the "shortest" interval, i.e. (0,1).

(^{21}) S's credence for ($\wedge_{i=1}^{n} a_i$) ought to be $\prod_{i=1}^{n} C(a_i)$ and her credence for ($\vee_{i=1}^{n} a_i$) ought to be $\sum_{I \subseteq \{1,...,n\}} (-1)^{|I|+1} \prod_{i \in I} C(a_i)$. But for any n, these will remain in (0, 1) so long as her credences for the individual a_is do.

(^{22}) The Lockean (about belief) is someone who thinks that believing p is just a matter of having a sufficiently high degree of belief for p (see, e.g. Foley (1992)). Let's say that the *Standard Lockean* is someone who thinks this about belief and wants to maintain that credences are

standard. This Standard Lockean might try to dig in his heels here and insist that suspension is not permissible in some of my cases since sufficiently high credence just is belief. My argument might be deployed as the *modus tollens* to his *modus ponens*: if credences are standard, then since suspension of judgment about *p* is epistemically permissible in these cases, believing *p* cannot just be a matter of having a (sufficiently) high standard credence for *p*. But this isn't just a stand-off. The relevant propositions are ordinary contingent propositions for which the subject has no relevant evidence, and as such, suspension of judgment is epistemically permissible. The Standard Lockean then owes us some story not only about why suspension about the relevant sort of *p* should be prohibited when one has no evidence either for or against *p*, but also one about why believing an ordinary contingent proposition should be permitted on absolutely no evidence.

([23]) See Hájek (2003) for a good discussion of probability gaps and their potential prevalence.

([24]) Some people have tried to argue that the belief-making threshold is shifty: that it can vary according to the circumstances of the subject, or different contexts of utterance. We can find suggestions like this in Hájek (1998), Weatherson (2005), Ganson (2008), Sturgeon (2008), Fantl and McGrath (2009) and a related one in Hawthorne and Bovens (1999). These accounts give up on the Straightforward Reduction strictly so-called since they say that belief and suspension will have to supervene on not just degrees of belief, but degrees of belief plus some other sorts of facts. Nonetheless, some of these views might be thought to respect the spirit of the Straightforward Reduction if not the letter, in that they can say that the only doxastic state that one needs to be in to believe or suspend is to have some relevant (standard) degree of belief. As such, I want to make clear that nothing shifty is going on in any of my cases. Standard mechanisms that make belief harder to come by are raised stakes or the making salient of the possibility of error. Absolutely nothing is at stake for the subjects in my cases. Even when stakes are as low as can be, subjects can rationally suspend with any degree of belief in (0,1) (assuming, of course, that credences are standard). And the same goes for the possibility of error. No mention of such a possibility is or need be made anywhere to make it that subjects can rationally suspend with very high and very low degrees of belief. Fix these other sorts of facts any way you like, we will still be able to get the result that the SJ-subinterval must be at least (0,1).

(^{25}) If S's credences for the cells of these partitions are each in $[0,1]$ but don't sum to 1, then those credences will fail to be countably additive or be in conflict with the normalization norm (or both).

(^{26}) I take it that this is fairly obvious. If a rational S's credences are uniformly distributed over these partitions, then it is very easy to see: her credence for each $p \in P_i$ will be $1/i$, and so as i goes up, her credences for each $p \in P_i$ will go down. But I have not demanded credal uniformity here. If rational S distributes her credences non-uniformly over P_i where i is very high, then not all of her credences for the propositions in P_i need be low. For instance, she may have a credence of 0.999 for some $p \in P_{10^{10}}$. But given that her credences for the propositions in $P_{10^{10}}$ ought to sum to 1, her credences for the remaining $10^{10} - 1$ propositions will have to sum to 0.001. Any way she disperses that 0.001 across those remaining propositions, her credences for those individual propositions will have to be very low. If i is sufficiently high, even if rational S's credence for some $p \in P_i$ needn't be extremely low, her credence for most of those propositions will have to be. In general, even without uniformity, for any P_i, a rational S's credence for at least one $p \in P_i$ will be at most $1/i$.

(^{27}) A uniform distribution over a countably infinite partition that is countably additive will violate the normalization axiom. Some (e.g. de Finetti (1970), famously) have worried about the claim that a rational subject must have a skewed credence function in these sort of cases. See Williamson (1999) for a good discussion.

(^{28}) Here is an informal and (hopefully) intuitive way to see this (adapted from Williamson (2007), p. 173; see Hájek (2003), p. 281–2 as well). First, I am assuming that the Straightforward Reductionist must avoid credence gaps for the reasons already discussed. Now, say Q is uncountable, and imagine a probability distribution (*Pr*) over Q. Now think of subsets of Q, Q_i as follows. Q_1 is the subset of Q whose members have probability $1/1$ or greater, Q_2 the subset of Q whose members have probability $1/2$ or greater, Q_3 the subset of Q whose members have probability $1/3$ or greater, and so on. For each i, Q_i has finitely many members. In particular, each can have at most i members or else the probability of the disjunction of its members would exceed 1. Now take any real number x such that $0 < x \leq 1$. If $Pr(q) = x$ then q is in at least one of the Q_is. For instance, if $Pr(q) = 0.001$, then q is in any subset of Q, Q_i such that $i \leq 1/1000$. And the same goes (*mutatis mutandis*, of course) for any x in $(0,1]$. But that

means that any $q \in Q$ that gets assigned a real number in (0,1] by Pr is Q_i for some i. Now take the union of these Q_is. This is a union of countably many finite sets. But a union of countably many finite sets is itself a countable set. This means that every $q \in Q$ such that $Pr(q) > 0$ is in this countable subset of Q. But Q is uncountable which leaves uncountably many $q \in Q$ that can only get probability 0. This means that a countably additive probability distribution over an uncountable partition will have to assign 0 to uncountably many cells of that partition. Given that we are assuming that S's credence function ought to be a (standard) probability function this means that S ought to have credence 0 for uncountably many $q \in Q$ (and credence 1 for their negations).

(^{29}) I am simply assuming that we can replace 'can know' or 'is in the position to know' in these cases with 'is epistemically permitted to believe' without incident.

(^{30}) Again, see Hawthorne (2002), and for some recent responses and discussion see Avnur (forthcoming) and Turri (forthcoming).

(^{31}) Almost. Someone like Tyler Burge in Burge (1993) might want to say that we can come to know ordinary contingent propositions a priori if we come to know them via certain kinds of testimony. Of course, my cases here do not involve testimony. See Christensen and Kornblith (1997) for a good discussion of Burge's view. We might also find a line of thought that can get us to the conclusion that there are a priori justified beliefs in ordinary contingent propositions in Douven (2008). These arguments rely on claims about our "epistemic goal" that seem to me highly contentious though. In fact, if anything, Douven's arguments make clear part of what is wrong with these teleological claims.

(^{32}) Although I have argued that belief and disbelief are not epistemically permitted in these cases, it is worth making clear that denying those claims does not secure the claim that suspension of judgment is not epistemically permissible. That claim would need an independent defence, whatever one said about the permissibility of belief and disbelief.

(^{33}) Although given our commonplace understanding of belief and degrees of belief these are quite bad results, there is obviously some space for a position that bites these bullets. This position could reduce suspending about p to having a standard credence for p (any at all), but claim that belief and disbelief are attitudes that a rational

subject has exactly when she commits to p or $\neg p$ in some distinct way (some way that doesn't at all involve degrees of belief). This would mean claiming (among other things) not just that believing p is something different from having credence 1 for p, but that it is rationally incompatible with having credence 1 for p. Obviously, we'd need to hear more about what belief is on this view before it could become a serious option.

(^{34}) See (e.g.) Lewis (1980) and Skyrms (1980) for the suggestion, and see Hájek (2003) for a good discussion of some of its drawbacks (both philosophical and mathematical).

(^{35}) Typically this is achieved by using sets of credence functions rather than a single function to represent a subject's total doxastic standing at a time. We can start with a set of credence functions R and the relevant set of propositions F that is the domain of the functions in R. We can then say that each $C(\cdot) \in R$ ought to be a probability function, and that a set of real numbers in the unit interval is assigned to the propositions in F. If any two functions in R disagree on their p-assignments (so that the resulting set is a non-singleton set), we can think of the resulting p-credence as imprecise (vague, mushy, etc.). It is also standard (although not required) to demand that R be convex and so demand that the sets assigned to the propositions in F be intervals.

(^{36}) For instance, Joyce (2005) claims that we should not try to capture states of "ambiguous or incomplete" (see p. 167 for more detail about what these amount to for him) evidence using a single credence function. He argues that picking any single credence function amounts to pretending to have information that one does not possess. He claims that when one's evidence is ambiguous or incomplete it is compatible with many distribution of objective probability over the hypotheses, so by distributing credences in any one way over them one ignores a vast number of possibilities that are consistent with one's evidence. We can find the suggestion that suspension of judgment be captured with credal imprecision in van Fraassen (1989) and Sturgeon (2010).

(^{37}) For some general worries about imprecise credences and the sorts of states they represent, see Elga (2010) and White (2010).

(^{38}) The canonical account of the rationality of this sort of chance-credence relation comes in Lewis (1980). Lewis proposes that the correct relation between chance and credence is given by his

Principal Principle, which states roughly that credences ought to track known chances.

(^{39}) In fact we can think of this as an extension of one of the oldest norms for suspension of judgment. In *Outlines of Pyrrhonism* Sextus says, "the term 'suspension' is derived from the fact of the mind being held up or 'suspended' so that it neither affirms nor denies anything owing to the equipollence of the matters in question" (Empiricus (1933), PH I, 196). When evidence supports all outcomes equally it is equipollent. Notice, the cases in which subjects have no evidence for or against any outcome are also cases in which they have equipollent evidence.

(^{40}) It is also worth making clear that genuinely making good on this line will leave the prospect of identifying suspension about *p* with, or reducing suspension about *p* to, having an imprecise *p*-credence extremely dim.

Access brought to you by: New York University