

Draft of 9/1/15; please contact me before citing.
sophie.horowitz@rice.edu

Accuracy and Educated Guesses

Belief, supposedly, “aims at the truth”. Whatever else this might mean, it’s at least clear that a belief has succeeded in this aim when it is true, and failed when it is false. That is, it’s obvious what a belief has to be like to get things right. But what about credences, or degrees of belief? Arguably, credences somehow aim at truth as well. They can be accurate or inaccurate, just like beliefs. But they can’t be true or false. So what makes credences more or less accurate? One of the central challenges to epistemologists who would like to think in degree-of-belief terms is to provide an answer to this question.

A number of answers to this question have been discussed in the literature. Some argue that accuracy, for credences, is not a matter of credences’ relation to what’s true and false, but to frequencies or objective chances.¹ Others are skeptical that there is *any* notion of accuracy can be usefully applied to credences, and that we should instead assess them according to their practical efficacy.² Yet another approach assesses accuracy using “scoring rules” – functions of the distance between credences and truth. According to this class of views, the closer your credence is to the truth (1 if the proposition is true, and 0 if it is false), the *better* it is, in a quite literal sense: scoring rules are understood as a special kind of utility function.³

This last approach – epistemic utility theory – has gained a significant amount of support in recent years. Part of its appeal is that it looks like a natural extension a common-sense thought about accuracy: that it’s *better* for our doxastic states to be right than wrong, and that for credences, it’s better to be close to the truth than far away.⁴ It is also a powerful bit of machinery, which can be used to justify or vindicate quite strong formal constraints on rational credence. But the approach faces problems as well. Just saying that close is better than far does not do much to narrow down the possible ways of

¹ For discussion these views, see Hájek [ms]. (Van Fraassen and Lange are among the defenders of frequentism; Hájek prefers objective chance.)

² See Gibbard [2008].

³ Supporters of this approach include Joyce, Greaves and Wallace, and Pettigrew, among others.

⁴ Joyce [2009] endorses this thought in the axiom he calls “Truth-Directedness”. Gibbard [2008] expresses the same idea in his “Condition T”.

measuring accuracy. And when we do try to narrow things down, defending the use of one scoring rule over another, we move farther and farther from the common-sense understanding of accuracy that we started with.

I won't enter this debate in depth here. Instead, I will propose a new way to understand accuracy. That is: we can evaluate credences' accuracy by looking at the "educated guesses" that they license. This framework is motivated by the thought that there is a straightforward way to assess credences' accuracy according to their relation to the truth, rather than to our practical aims – and by the common-sense thought that credences are more accurate as they get closer to the truth.

Here is the plan for the rest of the paper. In Section 1, I will introduce my proposal. In Section 2 I will argue that educated guesses can help us make sense of the phenomenon David Lewis calls "immodesty": the sense in which a rational agent's own doxastic states should come out looking best, by the lights of her way of evaluating truth-conduciveness or accuracy. (I will say much more about this in section 2.) As I'll argue, vindicating Immodesty is a minimum requirement for an account of accuracy, so it is good news for the guessing framework that it can be put to work in that way. In Section 3, I'll turn to the question of which formal requirements can be justified through this framework. I will argue that with some plausible constraints on rational guessing, we can use this framework to argue for probabilism; I will also briefly discuss some possible further applications, and alternative options for those who think that probabilism is too strong. In Section 4 I will (very) briefly survey two other accounts of accuracy, using them to bring out some of the strengths and weaknesses of the guessing framework.

1. Educated guesses

In gathering evidence and forming opinions about the world, we aim to get things right. If we're lucky, the evidence is decisive, and we can be sure of what's true and false. If we're unlucky – which is most of the time – the evidence is limited, and things are not so clear. In these cases, it's rational to adopt intermediate degrees of confidence, or credences. If we must act, we should do the best we can.

I want to look at a special kind of action: educated guessing. This is a type of action with the same correctness conditions as all-out belief. A guess is correct if it's true,

and incorrect if it's false. Guessing is something we are often called upon to do even when we're quite unsure what is the right answer to a question. As with any other action, if we must guess, it's rational to give it our best shot. The way to do that is to guess on the basis of our credences.

In short, guessing is a way that we can get things right or wrong, and rational guessing is done on the basis of our credences. In a relatively straightforward way, then, your credences can get things right or wrong by licensing true or false guesses. I'd like to propose that we make use of this connection to build an account of accuracy.

Specifically:

Your credences are *more* accurate insofar as they license true educated guesses. They are *less* accurate insofar as they license false educated guesses.

What are educated guesses? My characterization will be partially stipulative, but we won't end up too far from the everyday notion of guessing that we are all familiar with. To get an idea of the type of guesses I'm interested in, think of multiple choice tests, assertion under time or space constraints (such as telegrams), or statements like "if I had to guess, [P]... but I'm not sure..." More precisely, we can think of an educated guess as a potential forced choice between two (or more) propositions, made on the basis of your credences. If you are given some options – say, P and \sim P – and asked to choose between them, your educated guess should correspond to the option you take to have the best shot at being true.

Two important notes. First, the type of guesses I'm interested in are those that are licensed by your credences, and governed by rational norms. (I'll call them "educated" guesses to emphasize this.) Second, as I said before, guessing is an action, *not* a doxastic state. It is possible to rationally guess that P if you know or believe that P, or if you don't; in some cases, it may even be rational to guess that P if you rationally believe that \sim P.⁵ (See Question 2, below, for a possible example like this.)

⁵ One way in which my notion of guessing is somewhat stipulative is that, on my account, guessing that P is compatible with knowing that P. However, we would not normally describe acting on our knowledge as "guessing". Thanks to [OMITTED] for pointing this out.

What are the norms that govern educated guesses? As a start, here are three norms, which seem plausible enough to me (and which I'll assume for the rest of the paper):⁶

Simple questions: When faced with a forced choice between two propositions, your educated guess should be the proposition in which your credence is highest.

Suppositional questions: When faced with a forced choice between two propositions given some supposition, your educated guess should be the proposition in which your conditional credence (conditional on the supposition being true) is the highest.

Equal credence: With both suppositional and non-suppositional questions, if you have equal credence in both options, you are licensed to guess in favor of either one.

I'll be interested in the guesses that are licensed by a rational agent's credences, according to the norms above.

To get a handle on how these norms are meant to work, consider a couple of sample questions. Simple, non-suppositional questions are easy enough:

Q1: Is it raining?

In this case, if you are more confident of Rain than of \sim Rain, you're licensed to guess Rain. If you are more confident of \sim Rain, you're licensed to guess \sim Rain. If you are equally confident in both options, you may guess either way.

Suppositional questions are just slightly more complicated:

Q2: Supposing that it's not sunny, which is it: rain or snow?

⁶ At the moment I'll keep things simple and just look at two-option cases, but there is no reason I can see why the framework couldn't be extended to choices between three or more options. How the framework would develop, if expanded in this way, is an interesting question – it would likely turn out that, on any plausible expansion, licensed guessing will be partition-relative. Would that be a good thing, or a bad thing? Possibly, not so bad. See Lin and Kelly [2011] for an argument that partition-relativity is good – as applied to theory acceptance, rather than guessing. Similarly, Schaffer [2004] argues that knowledge is question-relative. Thanks to [OMITTED] and [OMITTED] for helpful discussion here. These points deserve further attention, but I will set them aside for present purposes.

Suppose your credences in these three (disjoint⁷) possibilities are as follows, where Cr is your credence function:

$$\text{Cr}(\text{Sun}) = .75$$

$$\text{Cr}(\text{Rain}) = .2$$

$$\text{Cr}(\text{Snow}) = .05$$

By your lights, then, it's most likely sunny. But Q2 asks you to suppose that it's *not* sunny. In response to this question, your credences license guessing Rain: given that it's not sunny, you regard it as more likely to be raining than snowing.

Your guesses can then be assessed straightforwardly for truth and falsity: either it's raining, or it's not. Suppositional guesses won't be assessed at all in cases where the supposition is false.

That's all I'll say for now about what guessing is, and when it's licensed. Does the guess framework give us a plausible account of accuracy? One way to test it is to see how well it fits together with the rest of our epistemological picture. I'll begin to explore this question in the next two sections.

2. Immodesty

In this section I'll argue that educated guesses can be used to vindicate "immodesty": roughly, the thesis that an epistemically rational agent should regard her own credences as giving her the best shot at the truth, compared to any other (particular) credences. The argument here will rely on the three norms for licensed guesses introduced in the last section. For this section, I will also assume probabilism: the thesis that rational credences are probabilistically coherent. (I will come back to probabilism in Section 3.)

What is immodesty, and why should we accept it? The term comes from David Lewis, who introduces it with the following example. Think about *Consumer Reports*, a magazine that ranks consumer products. Suppose that this month, *Consumer Reports* is ranking consumer magazines. What should it say? If *Consumer Reports* to be trusted, Lewis argues, it must *at least* recommend itself over other magazines with different product-ranking methods. Suppose *Consumer Reports* was "modest", and recommended

⁷ Pretend they are disjoint. As I'm writing this, it's sunny and raining at the same time.

Consumer Bulletin instead of recommending itself. Then its recommendations would be self-undermining, or inconsistent, in a problematic way. On p. 3, say, *Consumer Reports* recommends the Toasty Plus as the best toaster. On p. 7 it recommends *Consumer Bulletin*. Then, when you open up *Consumer Bulletin*, you find out that it recommends the Crispy Supreme. Which toaster should you buy? *Consumer Reports* is giving you incoherent advice. It can't be trusted.⁸

Lewis's example needs a few qualifications. Without saying more about the situation, it's not clear that *Consumer Reports* really should rank itself *best*. For instance, if *Consumer Bulletin* reviews a wider variety of products, it might be reasonable for *Consumer Reports* to recommend it as the best consumer magazine. It would also surely be reasonable for *Consumer Reports* to admit that *some* possible magazine could be better – say, *God's Omniscient Product Review Monthly* – especially if it does not have access to *GOPRM*'s testing methods or recommendations. What *Consumer Reports* can't do, on pain of incoherence, is recommend a magazine that (a) ranks the same products, (b) on the basis of the same information, but (c) comes out with different results.

Carried over to epistemology, the idea is that a rational agent should regard her own credences as optimal in the same sense as *Consumer Reports* should regard its own recommendations as optimal. Compared to other credences she might adopt – ranging over the same propositions, and on the basis of the same evidence – a rational agent should regard her own credences as giving her the best shot at the truth. To see why immodesty should be true for doxastic states, just imagine an agent who believes that it's raining, but also believes that the belief that it's *not* raining would be more accurate. This would be inconsistent and self-undermining – it would indicate that something has gone wrong, either with the agent's beliefs or with her way of assessing accuracy. The same should be true of credences: if credences are genuine doxastic states, aiming to represent the world as it is, they must aim at accuracy in the way that belief aims at truth. So if an agent has both rational credences and an acceptable way of assessing accuracy, she will be immodest.

⁸ See Lewis [1971]. Lewis defines "immodesty" slightly differently – in his terms an "inductive method", rather than the person who follows it, is immodest. (An "inductive method" can be understood as a function from evidence to doxastic states.) I'll follow Gibbard [2008] here in calling credences, or an agent who has those credences, immodest.

I understand immodesty as a kind of coherence between rational credences and the right account of accuracy. Given the right account of accuracy, credences that aren't immodest aren't rational; given rational credences, an account of accuracy that makes those credences modest isn't a good account.⁹ What I'll be doing here is arguing that, given the assumption that rational credences are probabilistically coherent, the guessing framework delivers immodesty. Since probabilism is a plausible and popular constraint on rational credence, I think this is a significant step in favor of the guessing framework. However, to show that guessing can do *everything* we want from an account of accuracy, we might also want to use it to argue *for* probabilism. I'll set this possibility aside until the next section.¹⁰

We are now ready to show how the guessing framework delivers Immodesty. This involves introducing a cleaned-up principle that expresses Immodesty in terms of educated guesses, and then showing why this principle is true.

First, here is the principle:

Immodesty: A rational agent should take her own credences to be best, by her current lights, for the purposes of *making true educated guesses*.

The guessing defense of Immodesty asks us to see epistemically rational agents as analogous to students preparing to take a multiple choice test. Even if you aren't sure of the right answers – after all, you don't know everything – you should take your best shot. Of course, we aren't *actually* preparing for a test like this, just as we aren't (usually) preparing to meet Dutch bookies or other potential money-pumpers. But imagining this scenario will help us show why Immodesty is true; it will help us show that insofar as you're rational, you take your credences to license the best guesses.¹¹

⁹ Joyce [2009] makes a similar claim about his principle, Admissibility, which claims that rational credences will never be weakly accuracy-dominated. (p. 267)

¹⁰ A final clarification about immodesty, before proceeding: immodesty is *not* a requirement that rational agents hold some particular attitude – for instance, that they know or believe that their credences are the most accurate. (Given some extra assumptions, we might argue that immodest agents have propositional justification for these things – but we don't need to get into that at the moment.) Agents can be immodest even if they have never considered questions about their own credences' accuracy; their credences must simply fit together with their notion of accuracy.

¹¹ My strategy here is directly based on the one employed by Gibbard [2008], discussed further in Section 4. Gibbard argues that we should assess our credences for their “guidance value”, or their ability to get us what we want, practically speaking. His argument, based on a proof by Schervish, involves imagining a

To see how Immodesty follows from the guessing picture, consider the following hypothetical scenario. You will take an exam. The exam will consist of just one question regarding a proposition (you don't know which one, beforehand) in which you have some degree of credence. You will have to give a categorical answer – for example, “It’s raining” – as opposed to expressing some intermediate degree of confidence. You will not have the option of refusing to answer. For the purposes of this exam, you only care about answering truly. Now suppose that you are choosing a credence function to take with you into the exam. You will use this credence function, together with the norms for guessing, to give answers on the exam. Which credence function should you choose? What we are interested in is which credence function does well by your current lights. So we will be considering various different candidate credence functions and evaluating their prospective success according to your *current* credence function. My claim is that if you are rational, then the prospectively best credence function, by your current lights, is your own.

For concreteness, let's call your *current* credence function “Cr”, and the credence function you *should* pick for the purposes of guessing “Pr”. So more precisely, my claim is that $Pr = Cr$. You should pick your own credences as the best credences to use for guessing.¹²

hypothetical series of bets. It might be helpful to think of my general line of argument as a “depragmatized” version of Gibbard’s. Gibbard points out that of course we aren’t really preparing for any such bets, and nor are we choosing our credences for that purpose – but it is “as if” we are. I want to take this stance towards my hypothetical quiz, as well. (Thanks to [OMITTED] for pressing me on this point.) The test scenario, as the bet scenario, shouldn’t be taken literally – it is still a useful illustration even if we know we won’t encounter the relevant bets. And we needn’t require agents to have beliefs or credences about which questions they’ll encounter, or to even consider potential guessing scenarios at all. (In fact, there are reasons to refrain from doing so, both for my strategy and for Gibbard’s: if there are an infinite number of potential questions, it’s impossible for agents to have positive credence, of each question, that that’s the question they’ll encounter. Thanks to [OMITTED] for pointing this out.)

¹² Some might object to the thought that there is just one credence function that you should pick, given your evidence. After all, if permissivism is true, many different credence functions are rational given your evidence. However, I don’t think that the current line of argument begs any questions against permissivism, at least if permissivism is understood *interpersonally*. Interpersonal permissivists should still accept immodesty – and indeed, may want to appeal to it as an explanation for why agents should not switch from one rational credence function to another without new evidence. See Schoenfield [2014] for an endorsement of immodesty in this context: Schoenfield argues that a rational agent should stick to *her* “epistemic standards” rather than switching because she should regard her own standards as the most truth-conducive.

To see how the argument works, we can start off by looking back at Q1 and Q2. (These will just be warmup questions; the real argument for Immodesty will come with Q3.) Suppose the exam question is Q1:

Q1: Is it raining?

Whatever credence function you choose for Pr will license guessing “yes” if $\text{Pr}(\text{Rain}) \geq .5$, and “no” if $\text{Pr}(\text{Rain}) \leq .5$. Suppose your credence in Rain is .8. Then, by your current lights, a “yes” answer has the (uniquely) best shot at being right. So you should pick a Pr such that $\text{Pr}(\text{Rain}) > .5$.

Simple questions like Q1 impose some constraints on Pr. In particular, Pr needs to have the same “valences” as Cr. That is, Pr needs to assign values that, for every proposition it ranges over, are on the same side of .5 as the values that Cr assigns. But questions like Q1 are not enough to fully prove Immodesty. To do well on Q1 and questions like it, you don’t need to pick Pr such that $\text{Pr} = \text{Cr}$. In this example, Pr could assign .8 to Rain, like Cr does, or it could assign .7 or .9. In fact, to do well on questions like Q1, you might as well round all of your credences to 0 or 1, and guess based on this maximally-opinionated counterpart of Cr.

More complicated questions impose stricter constraints on Pr. For example:

Q2: Supposing that it’s not sunny, which is it: rain or snow?

Suppose again that your credences in Sun, Rain, and Snow are as follows:

$$\text{Cr}(\text{Sun}) = .75$$

$$\text{Cr}(\text{Rain}) = .2$$

$$\text{Cr}(\text{Snow}) = .05$$

For this question, you need to be more picky about which credence function you choose for Pr. You will not do well, by your current lights, if you guess based on the maximally-opinionated counterpart of Cr. That credence function assigns 1 to Sun, and 0 to both Rain and Snow. So that credence function will recommend answering Q2 by flipping a coin or guessing arbitrarily. But, by your current lights, guessing arbitrarily on Q2 does not give you the best shot at guessing truly; it’s better to guess Rain. So you need to pick Pr such that it licenses guessing Rain, and does not license guessing anything else, on Q2.

To answer questions like Q2, then, you need to not only choose credences with the same valences as yours, but credences that also differentiate among unlikely

possibilities in the same way that Cr does. But this still does not show that $Pr = Cr$. You could do well on Q2, for example, by choosing a credence function that is uniformly just a bit more or less opinionated than Cr. This credence function is not Cr, but it will do just as well as Cr on questions like Q2.

Now consider another, more complicated question. For this example, suppose $Cr(\text{Rain}) = .8$.

Q3: A weighted coin has “Rain” written on one side, and “~Rain” on the other. It is weighted .7:.3 in favor of whichever of Rain or ~Rain is true. Now suppose:

- (a) the coin is flipped, out of sight;
- (b) you answer whether Rain; and
- (c) you and the coin disagree about Rain.

Who is right?

In this case, the best answer by the lights of Cr is that you are right. So you should choose a Pr that will also answer that you are right. I’ll first go through the example to show why this is, and then argue that questions like Q3 show that Immodesty is true.

We can work out why you should guess that you are right, in Q3, as follows. Since your credence in Rain is .8, you can work out that you will answer “Rain”. The only situation in which you will disagree with the coin, then, is one in which the coin lands “~Rain”. So we are comparing these two conditional credences: $Cr(\text{The coin is right} \mid \text{The coin says “~Rain”})$ and $Cr(\text{The coin is wrong} \mid \text{The coin says “~Rain”})$.

First, your credence that the coin will say “~Rain” is given by the following sum:

$$Cr(\text{The coin says ~Rain and it's right}) + Cr(\text{The coin says ~Rain and it's wrong})$$

Plugging in the numbers, using the weighting of the coin and the values that Cr assigns to Rain and ~Rain, we get: $(.7 * .2) + (.3 * .8) = .38$.

Your conditional credence that the coin is right, given that it says ~Rain, is $(.7 * .2) / .38 = .37$. Your conditional credence that the coin is wrong, given that it says ~Rain, is $(.3 * .8) / .38 = .63$. Since the second value is higher, the best answer by the lights of Cr is that, given that you disagree, you are right and the coin is wrong.

Questions like Q3 could be constructed with any proposition, and any weighting of the coin. To do well on the exam, when you don't know what question you will encounter, you need to be prepared for any question of this form. So you need to pick Pr such that it will give the best answers (by the lights of Cr) given any question like Q3 – involving any proposition and any possible coin.

The guesses that any credence function licenses on questions like Q3 depend on the relationship between the value that credence function assigns to the proposition (in this case, Rain) and the bias of the coin. If the credence function is more opinionated than the coin (in this case, if $Pr(\text{Rain}) > .7$), it will license guessing in favor of yourself. If the credence function is less opinionated than the coin (in this case, if $Pr(\text{Rain}) < .7$) it will license guessing in favor of the coin.

This is what we need to show that Immodesty is true. Suppose you choose a Pr that is different from Cr , so it assigns a different value to at least one proposition. Then, there would be at least one question for which Pr will license the “wrong” answer, by the lights of Cr . For example, suppose that $Cr(\text{Rain}) = .8$, but $Pr(\text{Rain}) = .6$. Then Pr will license the wrong answer in Q3: it will license guessing that the coin is right and you are wrong. This is because while Cr 's value for Rain is *more* opinionated than the weighting of the coin, Pr 's value for Rain is *less* opinionated. And it's easy to see how the point generalizes. To create an example like this for any proposition, P , to which Pr and Cr assign different values, just find a coin whose weighting falls between $Cr(P)$ and $Pr(P)$. Then, in a setup like Q3, Cr and Pr will recommend different answers. And by the lights of Cr , Pr 's answer will look bad; it won't give you the best shot at getting the truth.

To guarantee that Pr will license good guesses in every situation, Pr must not differ from Cr . So Immodesty is true: you should choose your own credence function, Cr , for the purpose of making educated guesses. $Pr = Cr$.¹³

¹³ Here is the more general form of Q3, and a more general explanation for why it delivers Immodesty:

Q3*: A weighted coin has P written on one side, and $\sim P$ on the other. It is weighted $x:1-x$ in favor of whichever of P or $\sim P$ is true, where $0 < x < 1$. Now suppose:

- (a) the coin is flipped, out of sight;
- (b) you answer whether Rain; and
- (c) you and the coin disagree about Rain.

Who is right?

Suppose $Cr(P) > Cr(\sim P)$; turn the example around if the opposite is true for you. You should guess in favor of yourself if $Cr(P) > x$, and in favor of the coin if $Cr(P) < x$.

2. Probabilism

We have now seen how educated guessing works, and how it delivers Immodesty. A rational agent should take her own credences to be the best guessers. This is a necessary condition on the right account of accuracy. But we might want more from accuracy: we might want to give accuracy-based defenses of certain rational coherence requirements. Since my defense of Immodesty assumed probabilism, we might hope that the guessing framework could be used to defend probabilism as well. The task is particularly pressing if we take educated guessing to be a rival of epistemic utility theory, which (usually) aims to deliver both probabilism and immodesty.

I'll argue in this section that we *can* use educated guesses to argue for probabilism. However, if readers find this argument contentious (as is inevitable: every existing argument for probabilism has its detractors) I hope they will still be interested in seeing what the guessing framework can do: either as a supplement to an independent argument for probabilism, or as a way to justify weaker coherence requirements like Dempster-Schafer. Section 3.4 offers some options along these lines.

Probabilism is traditionally expressed in three axioms. I'll use the formulations listed below. Assuming that \Pr is any rational credence function, T is a tautology, and Q and R are disjoint propositions, the axioms are:

Non-Triviality:	$\Pr(\sim T) < \Pr(T)$
Boundedness:	$\Pr(\sim T) \leq \Pr(Q) \leq \Pr(T)$

The probability that the coin says $\sim P$ will be the sum
 $\text{Cr}(\text{The coin is right} \mid \text{The coin says } \sim P) + \text{Cr}(\text{The coin is wrong} \mid \text{The coin says } \sim P)$

Or:

$$(1-y)(x) + (y)(1-x)$$

The following therefore gives you your conditional credences:

$$\begin{aligned} \text{Cr}(\text{Coin is right} \mid \text{Coin says } \sim P) &= (1-y)(x) / ((1-y)(x) + (y)(1-x)) \\ &= (x - xy) / ((1-y)(x) + (y)(1-x)) \\ \text{Cr}(\text{Coin is wrong} \mid \text{Coin says } \sim P) &= (y)(1-x) / ((1-y)(x) + (y)(1-x)) \\ &= (y - xy) / ((1-y)(x) + (y)(1-x)) \end{aligned}$$

To see which of the conditional credences will be higher, just look at the numerators (the denominators are the same). It's easy to see that if $x > y$, the first conditional credence will be higher than the second; if $y > x$, the second will be higher than the first. So you should guess that *the coin is right*, conditional on disagreeing, if your credence in P is greater than the weighting of the coin. You should guess that *you are right*, conditional on disagreeing, if your credence in P is less than the weighting of the coin.

Finite Additivity: $\Pr(Q \vee R) = \Pr(Q) + \Pr(R)$

My strategy here will be to show that if you violate Non-Triviality or Boundedness, you will either be guaranteed to guess falsely in situations where guessing falsely is not necessary, or you will miss out on a guaranteed-true guess in situations where it is possible to have one. Given some additional rational constraints on guessing, therefore, it is irrational to violate these axioms. I will then give a different kind of argument for Finite Additivity.

The rough idea behind my additional norms for rational guessing is as follows: it's irrational to guess falsely when it could be avoided. And it's irrational to fail to guess truly when you have the opportunity. (Alternatively, the *very* rough idea is: believe truth! avoid error!) Of course, to be plausible as rational norms, they must be spelled out further. Here is what I suggest:

No Self-Sabotage: Your credences are irrational if they uniquely license a guaranteed-false educated guess, in a situation where that could be avoided: that is, in a situation where you could adopt different credences, in response to the same evidence, that would *not* uniquely license that guaranteed-false guess.

No Missing Out: Your credences are irrational if they *fail* to license an educated guess that is guaranteed to be true in a situation where that could be avoided: that is, in a situation where you could adopt different credences, in response to the same evidence, that *would* license a guaranteed-true-guess.

I'd like to propose No Self-Sabotage and No Missing Out as *rational* norms. These norms place constraints on your credences by constraining the guesses that your credences can permissibly license. Their role is therefore a bit different from the first three norms, in Section 1, which describe how guessing is licensed on the basis of your credences. I'll argue that given these two norms, it is irrational for your credences to violate Non-Triviality and Boundedness.

Before putting these norms into action, a bit more about what they say, and how they are motivated. No Self-Sabotage is a prohibition on unnecessary, guaranteed-false guessing – being uniquely licensed to make a guaranteed-false guess *when doing so could be avoided*. Compare the following two situations. First, suppose you’re given a choice between two propositions that you’re certain are false. This is just a bad situation: both will be permitted (by Equal Credence), but you’ll be wrong either way. Since there’s no way out of making a false guess, guessing falsely – even making a guess that’s guaranteed to be false – shouldn’t be held against you. Second, suppose you’re given a choice between two propositions, and you’re *not* certain that both are false. But at least one of those propositions is guaranteed to be false – it’s a logical contradiction, say. In this second situation, your credences will uniquely license guessing in favor of one or the other. What No Self-Sabotage says is that in *this* kind of situation, something has gone wrong if your credences license you to make a guess that is guaranteed to be false.

No Missing Out says that your credences should license making guaranteed-true guesses whenever possible. Something has gone wrong, I propose, if you *could* be licensed to make a guaranteed-true guess – if other credences you could have *would* license a guaranteed-true guess, on the basis of the same evidence – but you’re not.

No Self-Sabotage and No Missing Out both include a provision that the agent’s evidence stay the same. This provision is important because of self-verifying cases like Jennifer Carr’s “Handstand” scenario.¹⁴ Carr imagines that you learn from your perfectly reliable yoga teacher that the objective chance of your successfully doing a handstand (which you’ll try in a minute) depends on your credence that you’ll be successful: in fact, whatever credence you adopt *will be* the objective chance of your succeeding. In this scenario you could guarantee yourself a true guess by becoming completely confident, either that you’ll fail or that you’ll succeed. That’s because your evidence is in part constituted by your credence in the relevant proposition; so, when you change your credence, you change your evidence as well. Carr argues – and I agree – that you’re not

¹⁴ Carr [ms].

rationally required to adopt extreme credences in this case, despite the fact that intermediate credences miss out on guaranteed perfect accuracy.

What *is* required in these cases is a tricky question, and answering it is a crucial task as we spell out the relationship between rational credence and accuracy. Nevertheless, I want to set that question aside for the moment. The same-evidence provision allows us to ignore cases like Carr's for the time being. We will only consider cases where changing your credence in P does not change your evidence about P.

In this section and the next, I will again use "Pr" to designate a rational credence function, and "Cr" to designate your current credence function without presupposing that those credences are rational.

2.1 Non-Triviality

With those new rational norms in hand, we're now ready to look at the first of the probability axioms.

Non-Triviality: $\Pr(\sim T) < \Pr(T)$

Non-Triviality says that your credence in a tautology, T, must be greater than your credence in its negation, $\sim T$. We can prove this axiom into two parts. First suppose that $\text{Cr}(\sim T) > \text{Cr}(T)$. This immediately leads to problems: if you were asked to guess whether T or $\sim T$, you would be licensed to guess $\sim T$. But T is a tautology, and therefore guaranteed to be true. So your guess is guaranteed to be false. And it is *unnecessarily* guaranteed to be false: if your credence in T were greater than your credence in $\sim T$, your guess would not be guaranteed to be false. Even stronger, in fact: it would be guaranteed to be true! Therefore if $\text{Cr}(\sim T) > \text{Cr}(T)$, you violate both No Self-Sabotage and No Missing Out.

Second, suppose that $\text{Cr}(T) = \text{Cr}(\sim T)$. If you were asked to guess whether T or $\sim T$, you would be licensed to answer either way. This means that you would be licensed to guess $\sim T$, which is guaranteed to be false. This guess is also *unnecessarily* guaranteed false: if your credence in T were greater than your credence in $\sim T$, you would not be licensed to guess $\sim T$ in this situation, so you would not be licensed to make a guaranteed-

false guess. If $Cr(T) = Cr(\sim T)$, you violate No Self-Sabotage. (You do not violate No Missing Out, however, since you *are* licensed to make a guaranteed-true guess that T.)

In both cases, violating Non-Triviality entails violating our new norms on rational guessing. The way to avoid violating these norms is to obey Non-Triviality. So given our two norms, Non-Triviality is a requirement on rational credence.

2.2 Boundedness

Boundedness:	$Pr(\sim T) \leq Pr(Q) \leq Pr(T)$
---------------------	------------------------------------

Boundedness says that it is irrational for you to be *more* confident of any proposition than you are of a necessary truth, and it is irrational for you to be *less* confident of any proposition than you are of the negation of a necessary falsehood. One way to read this axiom is as saying that, of all of the possible credences you could have, your credence in necessary truths must be highest – nothing can be higher! And your credence in necessary falsehoods must be lowest – nothing can be lower! If we add in a plausible assumption about what this means, we can prove Boundedness within the educated guess framework.

The assumption is this: there is a maximal (highest possible) degree of credence, and a minimal (lowest possible) degree of credence. I'll also assume a plausible consequence of this assumption in the guessing framework. First: if you have the maximal degree of credence in some proposition, A, you are *always* licensed to guess that A when A is one of your choices. That is, if you are asked to guess between A and A*, your credences always license guessing A. (If $Cr(A) = Cr(A^*)$, of course, you are licensed to guess either way by Equal Credence.) Second: if you have the minimal degree of credence in some proposition, B, you are *never* uniquely licensed to guess B. That is, if you are asked to guess between B and B*, you are only licensed to guess B if $Cr(B) = Cr(B^*)$.

For simplicity, let's assume that your credences satisfy Non-Triviality, which we have already argued for. So, $Cr(\sim T) < Cr(T)$. Assuming that there is a maximal credence and a minimal credence, we can normalize any agent's credences, assigning the value 1 to the maximal credence and the value 0 to the minimal credence. So, if $Cr(T)$ is maximal, $Cr(T) = 1$. If $Cr(\sim T)$ is minimal, $Cr(\sim T) = 0$.

First, let's prove that your credence in T should be maximal; that is, $\Pr(T) = 1$. Suppose that $\text{Cr}(T) < 1$. Then, I will argue, you violate *both* No Self-Sabotage and No Missing Out.

To show this, we can return to a question like Q3 from the last section. Suppose that you're "competing" against a weighted coin, biased in favor of the truth about T. The weighting of the coin, x , is such that $\text{Cr}(T) < x < 1$. (That is: the coin is weighted $x:1-x$, in favor of the truth about T, and it is *more* opinionated about T than you are.) Suppose that you and this coin disagree about whether T. Given that supposition, you will guess that the coin is right and you are wrong.

This violates No Self-Sabotage. In guessing that the coin is right, you are making a guaranteed-false guess. ("The coin is right", in this case, is equivalent to " $\sim T$ ".) It also violates No Missing Out. You are missing out on a guaranteed-true guess in favor of T. So you violate both additional norms. It is irrational for $\text{Cr}(T)$ to be non-maximal.

For the second part of Boundedness, we must prove that your credence in $\sim T$ should be minimal. So, $\Pr(\sim T) = 0$. Again, we can use a question like Q3. Suppose that your credence in $\sim T$ is .2. Consider the following question:

Q4: A weighted coin has some contingent proposition – you don't know which one, but call it "R" – on one side, and $\sim R$ on the other. It is weighted .9:.1 *against* whichever of R or $\sim R$ is true. Now suppose that the coin is flipped out of sight.

Which is right? The coin (however it landed), or $\sim T$?

Here we want to show that you will guess $\sim T$, which is guaranteed false.

In Q4, the coin is weighted heavily *against* the truth about R. You aren't told what R is; without any more information, your credence that the coin will be right should be .1. Your credence in $\sim T$ is .2. Although your credences in both propositions are quite low, your credence in $\sim T$ is still higher – so, you are licensed to guess $\sim T$. But $\sim T$ is guaranteed to be false. Your non-minimal credence in $\sim T$ is causing the problem here: if your credence in $\sim T$ was minimal, you would have been licensed to guess in favor of the

coin, which is not guaranteed to come up false. So you should have minimal credence in $\sim T$.¹⁵

Violating Boundedness also entails violating our two norms, No Self-Sabotage and No Missing Out. You could avoid these problems by adhering to Boundedness. So your credence in T should be maximal, and your credence in $\sim T$ should be minimal.¹⁶

2.3 Finite Additivity

While Non-Triviality and Boundedness provide constraints on our credences in necessary truths and falsehoods, Additivity says that our credences in contingent propositions should fit together with one another as follows:

Finite Additivity: $P(Q \vee R) = P(Q) + P(R)$

Contingent propositions are not themselves guaranteed to be true or false. So violating Additivity – while it may lead to some *irrational* guesses – will not necessarily lead to Self-Sabotage or Missing Out. That means that our two norms will not be enough to establish Additivity as a rational constraint. I will provide a different kind of argument for Additivity, and then address a potential objection.

¹⁵ Again, here is the general recipe for creating examples like this. Suppose your credence in $\sim T$ is z , where $0 < z < 1$, so z is not the minimal credence. Consider the following question:

Q4*: A weighted coin has some contingent proposition R on one side, and $\sim R$ on the other. It is weighted $1-x:x$ <i>against</i> whichever of R or $\sim R$ is true, where $0 < x < z$. Now suppose that the coin is flipped out of sight. Your question is: which is right? The coin (however it landed), or $\sim T$?
--

If you have minimal credence in $\sim T$, you will be licensed to guess in favor of the coin, no matter how it is weighted. You will only be licensed to guess $\sim T$ if the coin is weighted 1:0 against the truth about R – which is a *necessary* guaranteed-false guess, so not a mark of irrationality.

¹⁶ Note that the Boundedness principle I defend is weaker than the more general Boundedness principle that some other approaches aim to justify. The more general principle says that there *should be* an upper bound to your credences, rather than assuming from the outset that there *is* one. For instance, we can use Dutch Book Arguments to show that you should never have credence greater than 1: if you did, you would be licensed to make bets that guarantee you a loss.

This stronger Boundedness principle can't be defended on the guessing picture. However, I am not convinced that this should worry us. When we associate credences with dispositions to bet, we can make sense of what it means to have credence greater than 1; so, we need an argument showing that this is irrational. But if we associate credences with dispositions to guess, it's not clear what it is to have credence greater than 1. You can be licensed to *always guess that A*, but you can't be licensed to "*more-than-always*" guess that A .

The guessing picture therefore leaves us free to argue that credence greater than 1 is impossible – so no further argument for its irrationality is needed. Insofar as it *is* irrational to bet at odds that would seem to be sanctioned by more-than-maximal credence, this is a form of practical, not epistemic, irrationality.

Suppose you have the following credences in these two independent propositions, Q and R:

$$\text{Cr}(Q) = .3$$

$$\text{Cr}(R) = .4$$

Additivity says that, if you are rational, $\text{Cr}(Q \vee R) = .7$. My argument will bring out the fact that, if you violate Additivity, the way you guess regarding Q and R will differ depending on how the options are presented to you. (This is in line with the interpretation of the Dutch Book argument adopted by Skyrms, who draws on Ramsey: “If anyone’s mental condition violated [the probability axioms], his choice would depend on the precise form in which the options were offered him, which would be absurd.”¹⁷) The intuitive strategy will be to create two guessing scenarios regarding Q and R, and show that you will guess one way if you consider the disjunction, and another way if you consider whether one of Q and R is true, but they are presented separately. I’ll discuss the significance of this after going through the example.

As before, the argument for Additivity is broken into two cases. First, suppose that $\text{Cr}(Q \vee R) = .9$ (higher than the credence recommended by Additivity). Now consider the following question:

Q5a: Coin A has “yes” on one side, and “no” on the other. It is weighted .8:.2, in favor of “yes” if $(Q \vee R)$ is true and in favor of “no” if $(Q \vee R)$ is false.

Now suppose:

- (a) the coin is flipped out of sight, and
- (b) you guess whether $(Q \vee R)$. Say “yes” if you guess $(Q \vee R)$, and “no” if you guess $\sim(Q \vee R)$.

Interpret the coin’s “yes” or “no” as answering whether $(Q \vee R)$.

If you and Coin A disagree, who is right?

This question is again very similar to Q3. You and the coin are both answering whether the disjunction $(Q \vee R)$ is true, and your credence in $(Q \vee R)$ is more opinionated than the coin’s weighting. (Intuitively: from your perspective, the probability that you’re right

¹⁷ Skyrms [1987]; citation from Ramsey [1926], p. 41.

about $(Q \vee R)$ is .9, but the probability that the coin is right is only .8. So your conditional credence that you are right, given that you disagree, should be higher than your conditional credence that the coin is right, given that you disagree.) You should guess that, if you and Coin A disagree, you are right and the coin is wrong.¹⁸

Compare Q5a to the following question, again supposing that $Cr(Q) = .3$, $Cr(R) = .4$, and $Cr(Q \vee R) = .9$:

Q5b: Coin A has “yes” on one side, and “no” on the other. It is weighted .8:.2 in favor of “yes” ($Q \vee R$) is true and in favor of “no” ($Q \vee R$) is false. Coin B has “Q” on both sides. Coin C has “R” on both sides. Now suppose:

- (a) all three coins are flipped out of sight,
- (b) you guess “yes” or “no” in response to this question: *Did at least one of Coin B and Coin C land true-side-up?* and
- (c) You and Coin A disagree: either you said “yes” and the coin said “no”, or you said “no” and the coin said “yes”.

Interpret the coin’s “yes” or “no” as answering whether at least one of Coin B and Coin C landed true-side-up.

Between you and Coin A, who is right?

Your credence that at least one of Coin B and Coin C landed true-side-up should be .7: after all, your credence that Coin B landed true-side-up is .3, your credence that Coin C landed true-side-up is .4, and Q and R are independent. So from your perspective, the probability that you will be right is .7. The probability that the coin is right, however, is .8. So your conditional probability that you will be right, given that you disagree, is *less*

¹⁸ Plugging in the numbers: since your credence in $(Q \vee R)$ is .9, you will guess “yes”. So if you disagree, that means the coin must have landed “no”. We are therefore comparing the following two conditional probabilities: $Cr(\text{Coin A is right} \mid \text{Coin A says “no”})$ and $Cr(\text{Coin A is wrong} \mid \text{Coin 1 says “no”})$.

Your credence that Coin A says “no” is given by this sum:

$$Cr(\text{Coin A says “no” and it’s right}) + Cr(\text{Coin A says “no” and it’s wrong})$$

Plugging in the numbers, we get $(.8 * .1) + (.2 * .9) = .26$.

Your credence that Coin A says “no” and it’s right is $(.8 * .1)$. So your conditional credence that Coin A is right, given that it says “no”, is **.31**. Your credence that Coin A says “no” and it’s wrong is $(.2 * .9)$. So your conditional credence that Coin A is wrong, given that it says “no”, is **.69**.

So you should guess that, if you disagree, you are right and Coin A is wrong.

than your conditional probability that the coin will be right, given that you disagree. You should guess that if you disagree, Coin A will be right.¹⁹

This combination of guesses illustrates the inconsistency in your credences. In Q5a, you are licensed to guess that if you disagree with Coin A, *you* will be right. In Q5b, you are licensed to guess that if you disagree with Coin A, *the coin* will be right. But the only difference between Q5a and Q5b was in how your guess about Q and R was presented: as a disjunction in Q5a, and as separate guesses on Q and R in Q5b. So if you are rational, you should *not* answer differently in Q5a and Q5b.²⁰

We can create a parallel setup for the case where your credence in $(Q \vee R)$ is *lower* than the credence recommended by Additivity. All we need is a Coin A', whose weight is between your credence in $(Q \vee R)$ and the sum of your credence in Q and your credence in R. (For example, if your credence in $(Q \vee R)$ is .51, we could weight the coin

¹⁹ Plugging in the numbers again: Your credence in Q is .3, and your credence in R is .4. You know that Coin B will say "Q" and Coin C will say "R". So your credence that at least one of Coin B and Coin C will land true-side-up should be .7. You should guess "yes". If you disagree with Coin A, then, that means that Coin A must have said "no".

Your credence that Coin A says "no" is given by this sum:

$$\text{Cr}(\text{Coin A says "no" and it's right}) + \text{Cr}(\text{Coin A says "no" and it's wrong})$$

Plugging in the numbers, we get $((.8 * .1) + (.2 * .9)) = .26$.

In this question, when you disagree with Coin A, you are each answering the question of whether at least one of Coin B and Coin C landed true-side-up. Your credence that Coin A says "no" and is right about that question is $(.8 * .3)$. So your conditional credence that Coin A is right, given that it says "no", is **.92**. Your credence that Coin A says "no" and it's wrong about that question $(.2 * .7)$. So your conditional credence that Coin A is wrong, given that it says "no", is **.53**.

So you should guess that, if you disagree, the coin is right and you are wrong.

²⁰ Here is the general recipe for examples of this form. Suppose that $\text{Cr}(Q) = x$, $\text{Cr}(R) = y$, and $\text{Cr}(Q \vee R) = z$. Now, suppose $z > x + y$. Compare the following two questions:

Q5a*: Coin A has "yes" on one side, and "no" on the other. It is weighted $v:1-v$, where $x + y < v < z$, in favor of "yes" if $(Q \vee R)$ is true and in favor of "no" if $(Q \vee R)$ is false. Now suppose:

- (a) the coin is flipped out of sight, and
- (b) you guess whether $(Q \vee R)$. Say "yes" if you guess $(Q \vee R)$, and "no" if you guess $\sim(Q \vee R)$.

Interpret the coin's "yes" or "no" as answering whether $(Q \vee R)$.

If you and the coin disagree, who is right?

Q5b*: Coin A has "yes" on one side, and "no" on the other. It is weighted $v:1-v$, where $x + y < v < z$, in favor of "yes" if $(Q \vee R)$ is true and in favor of "no" if $(Q \vee R)$ is false. Coin B has Q on both sides. Coin C has R on both sides. Now suppose:

- (a) all three coins are flipped out of sight,
- (b) you guess "yes" or "no" in response to this question: *Did at least one of Coin B and Coin C land true-side-up?* and
- (c) You and Coin A disagree.

Between you and Coin A, who is right?

You will guess in favor of yourself in Q5a*, and in favor of the coin in Q5b*.

.6:2 in favor of “yes” if $(Q \vee R)$ is true, and in favor of “no” if $(Q \vee R)$ is false.) Again, you will guess inconsistently: you will guess in favor of the coin when you consider $(Q \vee R)$ presented as a disjunction, and you will guess in favor of yourself when you consider Q and R separately.

This is irrational. You have no basis for treating $Q5a$ and $Q5b$ (or their counterparts, with coin A') differently from one another. But if you violate Additivity, your credences require you to treat the two cases differently.

Here is another way we could put the point. Your guesses in questions $Q5a$ and $Q5b$ reflect how you regard the strength of your evidence about Q and R . In $Q5a$, guessing in favor of yourself, over the coin, makes sense because you consider your evidence to be a stronger indicator of whether Q or R is true than the coin is. From the perspective of your evidence, trusting the coin over your own guess is a positively bad idea; it gives you a *worse* shot at being right. Compare this to your guess in $Q6b$. From the perspective of your evidence, as characterized in $Q6b$, trusting your own guess over the coin is a positively bad idea. But if the relevant evidence – the evidence bearing on Q , and the evidence bearing on R – is the same, and you are judging its strength in comparison to the very same coin, it doesn't make sense to guess differently in the two cases. Your credences should not license both guesses simultaneously. The only rational option is to obey Additivity.

I'd like to close by addressing two worries you might have about this argument. First: you might think that providing a different kind of argument for Additivity from the kind we had for Boundedness and Non-Triviality is a weakness of the guessing picture. After all, popular defenses of probabilism – Dutch Book arguments and epistemic utility theory – argue for all three axioms in a unified way. The Dutch Book argument says that agents with incoherent credences will be licensed to take bets that guarantee a net loss of money (or utility, or whatever you're betting on). Epistemic utility theorists argue that incoherent credences are accuracy-dominated by coherent credences, or else that incoherent credences fail to maximize expected epistemic utility. On the guessing picture, however, the argument for Additivity is, in a way, weaker than the arguments for the

other two: it gives us an illustration of tension in your credences, rather than pointing to something positively bad that will result from that tension. Is this a problem?

I'd like to propose that we think of Additivity differently from the other axioms. The argument I gave was meant to show how, if your credences violate Additivity, you will fail to make sense by your own lights. How reliable you take yourself to be regarding Q and R depends on how you are asked about Q and R – how the very same guessing situation is presented to you. This is the same sort of argument we might make to show that it is irrational to believe that John a bachelor, but also believe that he's married. Neither of these particular beliefs is guaranteed to be false in virtue of your holding both of them. But you will have beliefs that don't make sense by your own lights – at least if you understand what it is to be a bachelor, and what it is to be married. We could make the same kind of argument in favor of other informal coherence constraints: for example, to show that it is irrational to believe both P and *my evidence supports* $\sim P$. There is a kind of incoherence involved in holding both beliefs, even if doing so does not lead to a straight-out contradiction. In both cases, we might not have the security of a decisive proof on our side. But that doesn't show that the rational requirements in question don't hold.

Of course, my argument for Additivity relied on some controversial assumptions. Most obviously, I relied on the thought that if you have evidence bearing on $(Q \vee R)$ as a disjunction, that very same evidence bears on both Q and R, separately. This leads us to a second worry: does my argument for Additivity assume what it's trying to prove? After all, claiming that Q5a and Q5b are “the same question, presented differently” might seem to beg the question against an opponent of Additivity. An advocate of Dempster-Schafer theory, for instance, might argue that it's possible to have evidence bearing on $(Q \vee R)$ as a disjunction that has no bearing on either Q or R individually. My argument would do little to persuade a fan of Dempster-Schafer to be a probabilist. So you might think this shows that the guessing account can't really provide a strong justification of probabilism.

I take this to count in favor of the guessing account. It can be used to make sense of, and argue for, the axioms of probability for those who are sympathetic to certain background assumptions. But it is also flexible enough that, were we to deny these assumptions, we would still be able to make use of the general framework. (See the next

subsection for some suggestions to this effect.) The guessing picture can therefore serve as a backdrop for some of the substantive debates in formal epistemology. And the particular argument I proposed for Additivity makes clear where the substantive assumptions come in to those debates.

2.4 Other applications

We've seen how the guessing picture can help us argue for the probability axioms as constraints on rational credence. One might wonder whether there are other constraints that it can justify: can it do more? Or, for those skeptical of Additivity, can it do *less*? A full exploration of these questions is beyond the scope of this paper (indeed, one of my hopes for this paper is to point towards these questions, rather than answering all of them). But here is a brief survey of some further questions we might use the guessing framework to answer.

(A) Setting aside formal requirements for the moment, the educated guessing framework is potentially useful for contexts in which we want to draw a connection between credences and various all-out epistemic notions. A salient example is *reliability*, which is typically understood as the propensity to get things right and wrong in all-out terms. One place this might come in handy is in thinking about “higher-order” evidence: evidence about your own rationality, what your evidence is, or what it supports. Many epistemologists find it plausible that this kind of evidence should influence what credences are rational for you to adopt. A natural explanation for this is that impairments in rationality often go along with impairments in reliability as well.²¹ I have also argued elsewhere for an explicit connection between credences and educated guesses, in the

²¹ The temptation to speak in all-out terms is clear in much of the literature on higher-order evidence. For example, see White [2009]’s “Calibration Rule”, which states: “If I draw the conclusion that P on the basis of any evidence E, my credence in P should equal my prior expected reliability with respect to P.” See also Elga [2007]’s (similar) formulation of the Equal Weight View: “Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right...” Though both White and Elga work in a degree-of-belief framework, they often slip into all-or-nothing terms to describe how higher-order evidence should work. The guessing picture could help to make this connection more precise.

interest of spelling out how higher-order evidence works.²² Guessing could be used in contexts like this one to allow us to use degreed and all-out notions at the same time.

(B) Back to formal constraints: we could use the guessing framework to argue for Regularity by endorsing stronger versions of No Self-Sabotage and No Missing Out. A stronger version of No Self-Sabotage might say that rational credences will never license guaranteed-false guesses, unless it is unavoidable (because all of one's options are guaranteed to be false). That would mean that it's irrational to have minimal credence in any contingent proposition: doing so would license you to make a guaranteed-false guess when your choice is between that guaranteed-false proposition and the contingent proposition. Similarly, a stronger version of No Missing Out might say that, whenever one could be licensed to make a guaranteed-true guess, one should *only* be licensed to make guaranteed-true guesses. That would mean that it's irrational to have maximal credence in any contingent proposition. (The weaker version of these norms – which don't entail Regularity – were strong enough for Non-Triviality and Boundedness. For now, I am only endorsing the weaker norms.)²³

(C) If we adopt a dominance avoidance principle, we can use the guessing framework to argue for the following two principles:

- (i) For all P and Q: if P entails Q, then $\Pr(P) \leq \Pr(Q)$.
- (ii) For all P, Q, and R: if P entails Q, and Q and R are mutually exclusive, then if $\Pr(Q) > \Pr(P)$, then $\Pr(Q \vee R) > \Pr(P \vee R)$.

The dominance-avoidance principle is as follows: it's irrational to hold some credences, if they license *strictly more* false guesses than other particular credences you could have had.²⁴

To prove (i): Suppose that P entails Q, but $\Pr(P) > \Pr(Q)$. Now suppose that you are asked to guess: P or Q? You will be licensed to guess that P. In the state of the world

²² See [OMITTED].

²³ Thanks to [OMITTED] and [OMITTED] for (separately) suggesting this to me.

²⁴ The argument in this subsection draws directly on work by Branden Fitelson and David McCarthy. (See their [ms].) They defend Dempster-Schafer axioms using a similar dominance-avoidance norm in their formal framework, which is a decision-theoretic picture using comparative confidence and all-out belief. Fitelson and McCarthy also prove that probabilism can't be defended using those tools alone. Thanks to [OMITTED] for extensive discussion and help on this section.

where P is false and Q is true, your guess will be false. You could have avoided guessing falsely in that state of the world if $Cr(P)$ had not been greater than $Cr(Q)$. Therefore, Cr is weakly dominated by another credence function – one which is just like it in every respect, except that it obeys (i).

To prove (ii): suppose that P entails Q , $Cr(Q) > Cr(P)$, but $Cr(Q \vee R) \leq Cr(P \vee R)$. Then suppose you're asked to guess: $(Q \vee R)$, or $(P \vee R)$? You will be licensed to guess $(P \vee R)$. There are four possible states of the world consistent with our supposition that P entails Q :

w1: $P, Q, \sim R$

w2: $\sim P, Q, \sim R$

w3: $\sim P, \sim Q, R$

w4: $\sim P, \sim Q, \sim R$

Your guess, $(P \vee R)$, will be true in w1 and w3, and false in w2 and w4. Compare this to another credence function, Cr' , which is just like Cr except that it obeys (ii). Cr' will guess $(Q \vee R)$, which is true in w1, w2, and w3, and false in w4. Therefore Cr' licenses strictly more *true* guesses than Cr – it licenses a guess that's true in w2, where Cr 's guess is false, and licenses the same guesses everywhere else. Cr is therefore weakly dominated by Cr' .

These two constraints, (i) and (ii), are especially interesting. Together with Non-Triviality, they are sufficient for Dempster-Schafer. So by adding this dominance avoidance norm, we could prove Dempster-Schafer using the guessing framework.

I want to stay neutral here on whether decision-theoretic reasoning is appropriate in this context, and hence whether rules like dominance-avoidance are the right way to think about rational guessing. So I do not want to either endorse or rule this particular application of the framework. However, it is an interesting possibility for further exploration, and one that is available to those who are sympathetic thinking about epistemic rationality in these terms. There may also be other ways to argue for constraints like Dempster-Schafer without thinking in decision-theoretic terms. I will leave that question open for now.

(D) Finally: it is interesting to note that, with the exception of Additivity and Immodesty, none of the norms I have argued for have relied on particular numerical values for our credences. (Notice that the original three norms about when guessing is licensed are put in terms of comparative confidence.) Therefore, much of what I have argued for could be adopted by someone who is skeptical of numerical-valued credences or subjective probabilities, and instead more interested in comparative confidence or plausibility.²⁵ Such a person could adapt my argument for Immodesty to defend a similar principle: rather than arguing that a rational agent should regard her own (numerical-valued) credence function as the best guesser, she could argue that a rational agent should regard her own plausibility ordering as the best guesser.

I won't pursue any of these applications in depth here. I mention them only to highlight the flexibility of the guessing framework, and the number of purposes that it could be used for. I take it to be a virtue of the guessing framework that it does not force our hand in a number of debates, such as whether to adopt Probabilism or Dempster-Schafer. Such debates should take place in our theory of rationality, not our theory of accuracy. But the arguments for or against various positions should be articulable in terms of accuracy. What I've shown here is that the guessing picture has promise as a framework in which those debates can play out.

4. Alternative approaches

Let's take stock. So far I have introduced my new framework and shown how it might be used to account for Immodesty and probabilism. In this section I will look very quickly at two alternatives to my proposal, each of which also offers a defense of these two requirements. Epistemic utility theory evaluates credences using special utility functions, or "scoring rules". Another strategy, which I'll call "the practical approach", does away with truth and looks instead at which actions are rationalized by an agent's credences. I will discuss these two approaches only briefly, to bring out some salient features of the educated guess picture in comparison to its competitors.

²⁵ Thanks to [OMITTED] and [OMITTED] for helpful discussion here.

4.1 Epistemic utility theory

Epistemic utility theory (EUT) starts off with what I referred to earlier as the common-sense notion of accuracy: the thought that credences are more accurate as they get closer to the truth.²⁶ According to EUT, epistemically rational agents should adopt the credences that maximize expected “epistemic utility”, much as decision theory understands practically rational agents as taking actions that maximize expected utility. Epistemic utility function, or “scoring rules”, are functions of credences’ closeness to the truth.

Many epistemic utility theorists aim to justify probabilism. Most canonically, Joyce (in his [1998] and [2009]) argues that incoherent credences are “accuracy-dominated” by coherent credences. There are two important premises needed for this argument to go through. One is the assumption that dominance-style reasoning is appropriate in this context. The other is the acceptance of certain axioms that narrow down the range of permissible measures of accuracy.

To deliver immodesty within the EUT framework, we must accept similar assumptions. First, EUT sees immodesty as a matter of maximizing expected epistemic utility from one’s own perspective (that is, as assessed by one’s own credences and scoring rule). This requires us to think about rationality in a decision-theoretic way, much as the earlier dominance assumption did. Second: immodesty enters the EUT framework as a constraint that narrows down the acceptable range of scoring rules. (The acceptable scoring rules are those that allow rational agents to regard their own credences as best – that is, as maximizing expected accuracy.) But in order to narrow down acceptable scoring rules in this way, we need to either accept certain strong axioms on acceptable accuracy measures, or else just build in immodesty – understood as expected-utility maximization – as its own very strong assumption from the start.

EUT is an interesting and powerful framework, and I don’t hope to argue definitively that the guessing framework is better. But I do think that the guessing framework has some important advantages, at least regarding the two assumptions that I’ve highlighted. Both assumptions have been questioned. To take a few examples: Selim Berker has argued against “consequentialist” or “teleological” reasoning in epistemology, and Jennifer Carr and (separately) Jason Konek and Ben Levinstein have argued against

²⁶ See, e.g., Joyce [1998] and [2009], Greaves and Wallace [2006], and Leitgeb and Pettigrew [2010].

simple decision-theoretic interpretations of the EUT machinery.²⁷ (Some, like Konek and Levinstein, argue for a subtle alternative interpretation of the rules in question; Carr argues for an alternative interpretation of the notion of epistemic utility. Others, like Michael Caie and Hilary Greaves, bite the bullet and accept strange results of decision-theoretic reasoning.²⁸) The guessing approach, however, allows us to avoid this challenge, because it does not build in or require any consequentialist assumptions. We are free to supplement the guessing picture with decision-theoretic or consequentialist reasoning, but we aren't forced to do so; we are also free to be non-consequentialists. All we need to say is that for any proposition, a rational agent should have the credence that gives her the best shot, given her evidence, at guessing truly on questions regarding that proposition. We can, as Selim Berker puts it, “respect the separateness of propositions.”²⁹

The second assumption – in particular, the specific axioms that EUT requires to deliver its strong results – has been questioned as well. For instance, Joyce's “Normality” axiom says roughly that equal distance from the truth must be evaluated equally for accuracy. Allan Gibbard objects to this Normality, arguing that someone should count as purely concerned with the truth, or purely concerned with accuracy, even if she values closeness to truth much more highly than distance from error. Patrick Maher gives a similar objection to Joyce's “Symmetry” axiom. He also argues that the whole EUT account of accuracy is implausible because it rules out the “Absolute Distance” measure, according to which the accuracy of one's credence is equal to the absolute value its distance from the truth.³⁰ According to Maher and Gibbard, the scoring rules that EUT works with are too far from our ordinary conception of accuracy; we should be suspicious, then, that EUT's arguments for probabilism and immodesty are really “purely alethic”.

The guessing framework largely avoids this problem as well. The account of accuracy it offers, including its basic rules for when guesses are licensed, is very simple and intuitive. As we saw, these intuitive pieces were all we needed to see that the guessing framework delivers immodesty. Probabilism, of course, required some stronger

²⁷ See Berker [2013 a] and [2013 b]; Carr [ms]; Konek and Levinstein [ms].

²⁸ See Caie [forthcoming] and Greaves [forthcoming].

²⁹ Berker [2013b].

³⁰ See Maher [2003] and Gibbard [2008].

norms on rational guessing, which may be challenged. But the guessing framework allows this debate to take place in a simple and intuitive setting – norms regarding when guessing is rational and irrational.

4.2 The practical approach

An alternative family of arguments tries to justify rational requirements such as probabilism and immodesty by looking at *practical* value. These arguments don't appeal directly to any particular understanding of accuracy, or any other way of evaluating credences directly in relation to truth. Instead, the practical approach builds on the connection between credences and rational action, as understood by decision theory. The practical argument for probabilism is the Dutch Book Argument (DBA). If you have incoherent credences, the DBA says, you will be licensed to accept a series of bets that, together, guarantee a sure loss (of money, utility, of whatever you're betting on).³¹ The practical argument for immodesty, given by Gibbard in his [2008], involves imagining a continuum of bets at various odds. Your credences will license taking some bets, and rejecting others. Gibbard argues that you should take your own credences to be best for the purposes of betting, or acting more generally – other credences will recommend taking bets that, by your current lights, look bad, or rejecting bets that look good.

Practical arguments provide an economical way of accounting for requirements of epistemic rationality. They don't require us to posit a special kind of epistemic utility; instead, they piggyback on practical utility, which has independent uses in the theory of practical rationality. But there is reason to think that we should try to do better. Most obviously, the phenomena that these practical arguments attempt to explain are, at face value, purely epistemic. Why should *epistemic* rationality be held hostage to practical concerns, such as how much money you're likely to make? (We don't generally think that you should adopt one belief over another because of monetary gain – so how are these arguments different?) For those who want to maintain that the practical and the epistemic as distinct normative realms, practical arguments for epistemic requirements miss the mark.

³¹ The Dutch Book Argument originates in Ramsey [1926]. See Vineberg [2011] for an overview.

There is much more to be said here. For instance, defenders of “depragmatize” Dutch Book arguments interpret them as manifestations of an epistemic phenomenon, rather than taking on the practical aspects at face value.³² I will come back to this in the last section. (It’s worth noting, though, that not all defenders of the practical approach want to adopt this sort of understanding. Gibbard explicitly abandons hope for of a purely epistemic argument for immodesty.) But once again, the guessing approach allows us to avoid this challenge. The guessing arguments for immodesty and probabilism don’t need to be “depragmatized” – they are already non-pragmatic. For those who find the practical strategy unsatisfying, the guessing framework offers an improvement. It allows us to say that a rational agent should be coherent and immodest, not because this will make her happy or rich, but because she takes her credences to give her the best shot at representing the world as it is.

5. Some final comments on guessing and accuracy

In this last section I’ll return to our original question: what makes credences more or less accurate? Does the guessing framework give us the kind of answer we were looking for?

We started with the intuitive thought that credences are more accurate as they get closer to the truth – for example, it’s better to have .9 credence in a true proposition than .8. It may not be immediately obvious how the guessing picture does justice to this thought. After all, if you and I are both asked to guess whether P, your true guess is as good as mine – regardless of whether one of our credences is much *closer* to the truth. One might object that this is the wrong result. We should be able to explain why your .9 credence is more accurate than my .8.

The objector gets one thing right: the guessing picture doesn’t allow us to differentiate credences of .9, .8, and .50001 *if we only look at one guess*, the one between P and \sim P. But we can do justice to the original thought that “closer is better” if we look at *all* of the guesses that our credences license. Someone with .9 credence in a true proposition P will guess correctly not just about P, but about lots of *other* questions as well – questions which someone with .50001 credence will often get wrong. Think back to the continuum of weighted coins we imagined in the argument for Immodesty. You

³² See, for instance, Christensen [1996] and Howson and Urbach [1993].

should expect to “beat” a given coin if you’re more opinionated than the coin. So as your credence gets closer to a proposition’s truth-value, the space of possible coins that are better guessers than you gets smaller and smaller, and the space of possible coins that are worse guessers gets bigger and bigger. Understanding the coins as representatives for all of the possible guesses the your credences license – better than coin A, worse than coin B, etc. – we can see that in general, the space of possible questions you can expect to answer correctly gets larger and larger as your credence gets closer to the truth. Greater accuracy, on the guessing account, amounts to getting more and more true guesses.

Another important aspect of our everyday conception of accuracy is that it is an alethic notion. In this respect, the guessing framework captures our notion of accuracy much better than the practical picture. It is decidedly less pragmatic than Gibbard’s (explicitly pragmatic) “guidance” account. Unlike Gibbard’s story, the guessing account essentially involves the connection between credences and truth. It is also less pragmatic than the simple, straightforward interpretation of the Dutch Book account: it appeals to the desire for truth, rather than utility or money.

But is the guessing picture *completely* free from pragmatic concerns? Guessing is, after all, an action. And in any real exam, whether it’s rational to guess one way or another is going to be subject to all kinds of practical concerns. This raises the worry that the guessing account isn’t purely alethic after all.

I claimed earlier that the guessing arguments, unlike Dutch Book arguments, do not need to be “depragmatized”. But depragmatized Dutch Book arguments nevertheless give us useful guidance for how the guessing arguments should be interpreted. What we’re interested in isn’t the all-things-considered rationality of guessing, or preparing to guess – after all, guessing falsely might be all-things-considered rational in some cases, and in other cases we might know that we won’t have to guess at all. Rather, we should look at the guesses sanctioned by our credences, and see them as illustrations of underlying properties of those credences.

In fact, this aspect of the depragmatization strategy seems to me to work better for the guessing account than for Dutch Book arguments. For the guessing account, we need only look at one particular kind of action (answering whether propositions are true or

false), and one desire (to answer truly), to illustrate the epistemic phenomena we are interested in. This action and desire are much more directly connected to epistemic concerns than betting behavior is. As I mentioned before, guessing is already an action with familiar correctness conditions, which are the same as those of full belief. It is natural to think, therefore, that credences that license true guesses are better, epistemically speaking, than credences that license false ones. And if your credences unnecessarily license guaranteed-false guesses – or if your credences are more likely, by your own lights, to license false guesses than other credences you could have – it is irrational to hold those credences.

6. Conclusion

We now have a new formal notion: educated guessing. I have argued that it gives us a natural and plausible way to think about accuracy for credences. The framework is simple and non-committal, and it fits together with many other things we might independently want from our theory of epistemic rationality. It vindicates Immodesty, and with a couple of plausible norms, can be used to argue for probabilism as well. In addition, I have pointed to a number of other potential applications that the framework might be used for, and highlighted some advantages that it has over other popular approaches. I hope to have established the guessing framework as a valuable new addition to our epistemic picture: one that can do much of the work that we need, and which has some advantages over its competitors.

Alongside these advantages, the guessing account also raises some questions. How do educated guesses fit in with a plausible philosophy of mind? Are there implications for other epistemological issues, like knowledge or full belief? Morality? Practical rationality? Do the norms for rational guessing give us strong enough constraints on rational credence? And so on. As with any new philosophical tool, we can begin to answer these questions as we see what the tool can be used for. It is my hope that educated guesses can do quite a lot.

Bibliography

- Arntzenius, Frank. [2008] “Rationality and Self-Confidence”. In Tamar Szabó Gendler and John Hawthorne, ed. *Oxford Studies in Epistemology*, vol. 2. Oxford: Oxford University Press, pp. 165-178.
- Berker, Selim. [2013a] “The Rejection of Epistemic Consequentialism”. *Philosophical Issues* 23, pp. 363-387.
- [2013b] “Epistemic Teleology and the Separateness of Propositions”. *Philosophical Review* 122, pp. 337-393.
- Caie, Michael. “Rational Probabilistic Incoherence”. Forthcoming in *Philosophical Review*.
- Carr, Jennifer. [ms] “Epistemic Utility Theory and the Aim of Belief”.
- Christensen, David. [1996]. “Dutch-Book Arguments Depragmatized: Epistemic Consistency For Partial Believers”, *The Journal of Philosophy*, 93: 450–479.
- Elga, Adam. [2007] “Reflection and Disagreement”. *Noûs*, volume XLI, no. 3, pp. 478-502.
- Fitelson, Branden and McCarthy, David. [ms] “Toward an Epistemic Foundation for Comparative Confidence”. Available at: http://fitelson.org/cc_handout.pdf.
- Gibbard, Allen. [2008] “Rational Credence and the Value of Truth”. In Tamar Szabó Gendler and John Hawthorne, ed. *Oxford Studies in Epistemology*, vol. 2. Oxford: Oxford University Press. pp. 143-164.
- Greaves, Hilary. “Epistemic Decision Theory”. Forthcoming in *Mind*.
- Greaves, Hilary and Wallace, David. [2006] “Justifying conditionalization: Conditionalization maximizes expected epistemic utility”. *Mind* 115 (459): 607-632.
- Hájek, Alan. [ms] “A Puzzle About Partial Belief”.
- Howson, Colin, and Peter Urbach (1993), *Scientific Reasoning: The Bayesian Approach*, second edition, La Salle, Illinois: Open Court.
- Joyce, James. [1998] “A nonpragmatic vindication of probabilism”. *Philosophy of Science* 65 (4): 575-603.

- [2009] “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief”. In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*. *Synthese*.
- Kelly, Thomas. [2013] “How to be an Epistemic Permissivist”. In Ernest Sosa, Matthias Steup and John Turri (eds.), *Contemporary Debates in Epistemology*, 2nd edition. Wiley-Blackwell.
- Leitgeb, Hannes and Pettigrew, Richard. [2010] “An Objective Justification of Bayesianism I: Measuring Inaccuracy”. *Philosophy of Science* 77 (2): 201-235.
- Lewis, David. [1971] “Immodest Inductive Methods”. *Philosophy of Science* 38 (1): 54-63.
- Lin, Hanti and Kelly, Kevin. [2011] “A geo-logical solution to the lottery paradox, with applications to conditional logic”. *Synthese* 186 (2): 531-575.
- Maher, Patrick. [2002] “Joyce’s Argument for Probabilism”. *Philosophy of Science* 69 (1): 73-81
- Ramsey, Frank. [1926] “Truth and Probability”, in Henry E. Kyburg and Howard E.K Smokler (eds.), *Studies in Subjective Probability*, Huntington, NY: Robert E. Kreiger Publishing Co.
- Schaffer, Jonathan. [2004] “From Contextualism to Contrastivism”. *Philosophical Studies* 119:73–103.
- Schoenfield, Miriam. [2014] “Permission to Believe: Why Permissivism is True and What it Tells Us about Irrelevant Influences on Belief.” *Noûs* 48(2): 193-218
- Skyrms, Brian. [1987] “Dynamic Coherence and Probability Kinematics”, *Philosophy of Science*, 54: 1–20.
- Vineberg, Susan, "Dutch Book Arguments", *The Stanford Encyclopedia of Philosophy* (Summer 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2011/entries/dutch-book/>.
- White, Roger. [2009] “On Treating Oneself and Others as Thermometers”. *Episteme* 6 (3): 233-250.