Sophie Horowitz shorowitz@umass.edu

MEW 2026

What follows are two chapters from *Guesswork* (forthcoming with Princeton University Press). These cover the most important parts of the project, and everything I will present at MEW. This introduction contextualizes the project and tells you a little bit about the rest of the book. I am attaching both chapters in full, but you could just look at sections 1 and 2 of Chapter 3 and get everything you need to understand Chapter 4, where the actual arguments are.

Guesswork offers a theory of **accuracy** for **credences**, or degrees of belief. While many epistemologists (and social scientists) find it helpful to theorize with degrees of belief, it is somewhat mysterious in what sense such states can get things *right and wrong*. Unlike simple, all-or-nothing beliefs, intermediate belief states cannot be true or false. On the other hand, it does seem that the notion of accuracy applies to degreed states: if P is true, someone who is fully certain of P should count as perfectly accurate, and if one person is more confident of P than another, the more confident person should count as more accurate (in that respect). Can we say anything more than that?

This question is important to answer on its own terms. If we want credences to fit into the rest of our epistemology and philosophy of mind in much the same way that all-or-nothing beliefs would fit in, we need to explain how they can get things right or wrong. It is also important to answer because accuracy plays an important role in most theories of rationality and justification. Some theories, like Reliabilism, hold that justified beliefs are *objectively* likely to be true. For other theories, having a justified belief involves *believing* that conditions obtain which make one's belief likely to be true.

Over the last 25 years or so. a prominent branch of formal epistemology has emerged, giving accuracy a central role. "Epistemic utility theory" sees accuracy as a kind of utility, and degrees of belief can have more or less of it as they get closer to or farther from the truth. The accuracy of a person's belief state can be characterized in precise numerical terms, and the relationship between accuracy and rational belief is decision-theoretic. Using EUT, epistemologists have proposed accuracy-based arguments for many of the popular formal constraints on rational credence, such as probabilism, conditionalization, the Principal Principle, and the Principle of Indifference. EUT has been, by far, the most popular way to theorize about the accuracy of degrees of belief.

But while EUT has been enormously fruitful and popular, some epistemologists worry about its underpinnings. Berker ([2013a], [2013b]) has objected the consequentialist (or "teleological") nature of epistemic utility theory, and how it treats rational belief as a gamble where losses in one bet can be compensated by gains in another. Maher [2003] and Gibbard [2008] each raise a different sort of objection, arguing that the conception of epistemic value used in EUT theory goes far beyond accuracy, and so arguments in EUT should not be understood as accuracy-based. Furthermore, the work in EUT tends to be (excessively?) mathematical and can appear, sometimes rightly, to address different questions from those of traditional epistemology. Although there is more to say about all of these worries, they give us reason to ask whether there is another way.

Guesswork develops a new theory of accuracy, aiming to avoid as much as possible the controversial assumptions and heavy formalism that have raised objections to EUT. My goal is to see how far we can get with a more intuitive, and less committal, notion of accuracy. The core idea of my account is that degrees of belief correspond to all-or-nothing guesses in response to questions. For example, suppose you are 75% confident that the path home is to the left, and 25% confident that the path home is to the right. If you are asked to guess whether the path home is to the left or to the right, you should guess "Left". If that is the correct path, your guess will be true. This fact counts in favor of your accuracy. Your credences are *accurate*, I'll argue, insofar as they license true guesses. They are *inaccurate* insofar as they license false guesses.

The guessing account of accuracy is very simple, but quite powerful. In the book, I argue that we can use it to make sense of many popular formal or structural constraints on rational credence, as well as concepts like reliability, which are more standardly used in conjunction with on/off beliefs than with credences. Along the way, I explore the broader question of what it means for an argument to be "accuracy-based", and how different arguments appeal to different understandings of the connection between rationality and truth.

The first two chapters of the book provide an introduction and an overview of epistemic utility theory, setting up the challenges for EUT and the similarities and differences between the two projects.

The third and fourth chapters are attached here. Chapter 3 explains the main idea of the approach: what guessing is, and when and how your guesses count for and against your accuracy. Chapter 4 makes the case for the guessing account as a legitimate way to assess accuracy, by arguing that it passes two tests. First, it can explain *Truth-Directedness*: the fact that our credences become more accurate as they approach 1 in truths and 0 in falsehoods. Second, it can explain *Immodesty*: the fact that a rational agent will take her own credences to have the best prospects for accuracy, as compared to other credences she might adopt instead.

Chapters 5 and 6 explore coherence requirements, including some basic axioms of formal coherence. Chapter 6 looks at possibilities for arguing for probabilism. Chapter 7 discusses conditionalization, and Chapter 8 applies the guessing approach to "mushy" or imprecise credences. Chapters 9 and 10 discuss Regularity and the Principal Principal. Finally, in Chapter 11, I survey a few ways in which the guessing framework might prove useful outside of formal epistemology.

Chapter 3. Educated Guesses

Our central question is how we should understand accuracy for credences. The previous chapter looked at epistemic utility theory, the dominant approach in contemporary formal epistemology. This approach started with the premise that credences are more accurate as they get closer to the truth. *Perfectly* accurate credences would assign 1 to all truths and 0 to all falsehoods. And beyond that, it's more accurate to have higher credence in truths, and lower credence in falsehoods. If it's raining, credence 0.6 in rain is more accurate than credence 0.5, which is more accurate than credence 0.4, and so on. Epistemic utility theory takes it as axiomatic that closer is better.

On the approach I am pursuing here, we will recover a version of that thought. But that is not where we will start. Instead, we will look at a different relationship between credences and truth. I will argue that credences license *guesses*: given a forced choice between two propositions and the aim of choosing a true proposition, if possible, your credences sanction choosing one or the other. This proposition is your guess, and it is either true or false. We can understand the accuracy of your credences in terms of the extent to which they get things right or wrong, by licensing true or false guesses.[1]

1. Guessing: the basics

The basic idea of this approach is: credences are accurate insofar as they license true guesses, and inaccurate insofar as they license false guesses. We will have to define some terms to make sense of this idea. Then we will complicate it a little bit.

First, what are guesses? What I mean by guessing is something very close to at least one version of our everyday notion. A guess is something you'd assert, write, or think in response to a question, where only all-or-nothing answers are acceptable. A multiple choice test is one circumstance which calls for a guess in response. We guess when we're not sure of the answer to a question, but have to choose. I will understand guesses, here, as *actions*, but that action may well be mental. (We will also return, in Chapter 11, to the question of whether guessing is equivalent to full belief.)

Sometimes we talk about guessing as an alternative to asserting, believing, or knowing in order to emphasize that the guesser is in a less-than-ideal epistemic state: you guess on the final exam because you don't know the answer, or you say "I guess" to emphasize your uncertainty. Some philosophers have even argued that guessing is arational: for example, L. Jonathan Cohen writes that guessing is "a substitute for proper reasoning" rather than something which itself can be justified.[2] He goes on to argue that we can sometimes say we are guessing in order to signal that we do not stand behind our claim, and to abdicate responsibility for it. Similarly, Gendin [1998] argues that if one guesses that P, one "weakly endorses P", but also "considers his evidence for P inadequate to support the claim that he knows P."[3] After all, why would you guess P if you knew P? Another body of literature in the 1980s took up the question of whether

---

[1] This chapter draws most heavily on Horowitz [2019]. Parts of section 3 also draw on Horowitz [2021].

[2] Cohen [1974], p. 196.

[3] Gendin [1998], p. 436. Gendin goes on to discuss guessing in the context of so-called evidence for clairvoyance, a topic which I will unfortunately not take up here.

guessing is compatible with knowing, considering, for example, cases where someone answers exam questions reliably without realizing that they are doing so.[4] Sosa [2015] characterizes (mere) guessing as a kind of affirmation which aims at aptness (that is, truth), but which lacks various kinds of higher-order endorsement; he also distinguishes guessing from belief and judgment, though he writes that in some cases guesses can constitute knowledge.[5]

As I'll use the term, guessing *is* compatible with knowing, and with being sure.[6] It is compatible with having or lacking various different higher-order attitudes about one's own rationality or reliability; one need not see oneself as "merely guessing" to guess in my sense. It is also compatible with rationality, and subject to rational norms. I think it is natural to think of guessing as rationally evaluable in many contexts: we often talk about "good guesses" and "educated guesses", and one could easily imagine criticizing someone for choosing a multiple-choice answer that was less-supported by her evidence than one of the alternatives. As Sorensen [1984] points out, "Your guess is as good as mine" is not a tautology.[7] But to the extent that you think Cohen and Gendin are right, and that guessing implies a lack of epistemic standing — or to the extent that you agree with Sosa that guessing implies a lack of higher-order endorsement — you can take my account to be stipulative.

If we agree that guessing is subject to rational norms, we might find it plausible that rationality forbids us from guessing arbitrarily, and also forbids us from making guesses which we think are probably false. But on my account, we will often be licensed to make guesses of both types. This is because of a second stipulation: the guesses I am interested in are guesses in response to *forced choices*. This means, among other things, that if you are asked to choose between two options which you think are probably false, you will end up making a guess that you think is probably false. And in cases where neither option appears to you to be more likely than the other, you will end up guessing arbitrarily.

We can look at a couple of examples to illustrate these last two points. First, suppose someone asks you to choose between the following options:

> {The moon is made of green cheese; the next seven cars I see out this window will all be green Toyotas.}

Both of these options are implausible, and (I'm assuming) you take them both to be almost certainly false. But if you have to choose, I'm betting you'll choose the latter. Under the circumstances, this is the choice that makes sense — even, probably, rationally required. This case is one in which you will probably say (if you're answering out loud) something like, "I guess," or "if I have to choose…" to signal the fact that you think you probably *won't* see seven green Toyotas in a row. But you could also be offered a choice like this:

> {The moon is made of green cheese; an apple is a fruit.}

Here you'll pick the second option without hedging, since it is obviously true. (If you have any doubt, it will likely be due to the strangeness of the question; you might wonder whether your

---

[4] See, for example, Sorensen [1982] and Dartnall [1984], among others, which respond to Colin Radford's earlier paper "Knowledge — by Examples" (Radford [1966]).
[5] See Carter [2020] for further critical discussion of this view.
[6] Here I agree with Cohen [1974], p. 195. Although he argues that guessing is arational, he also takes guessing P to be compatible with a range of doxastic attitudes towards P, including belief, suspicion, and certainty.
[7] Sorensen [1984], p. 78.

interlocutor is trying to trick you.) This illustrates a couple of facts about guessing, in the sense I'm interested in: that you can reasonably guess regardless of whether it would make sense to *say* "I guess", and that you can reasonably guess regardless of whether you think your guess is likely to be true.[8]

This brings us to one more stipulation: we will abstract away from contextual factors that might influence the way you respond to questions in conversation. When someone poses a forced choice question in conversation, they may implicate all sorts of things, and impart all sorts of information just by their choice of options or their framing of the question. Similarly, your own goals or values might prompt you to give one answer even though you take the other to be more likely correct: if you don't want to hurt someone's feelings, for example, or you don't want to ruin their birthday surprise. Recent work by Dorst and Mandelkern argues that in certain contexts (in particular, where there are more than two available options) one might make more or less specific guesses, depending on whether one merely wants to be accurate, or also wants to be informative.[9] For instance, asked where a certain applicant will enroll in law school — where the possible law schools are Yale, Harvard, Stanford, or NYU — one might permissibly answer "Yale" or "Yale or Harvard". ("Yale" is more informative, since it rules out more possibilities, but "Yale or Harvard" is more likely to be true. Other answers might be permissible as well given the details of the case.)

We will set aside all of these complicating factors here, and focus on a particular type of context: one in which your only aim is to guess truly, and in which no further information is conveyed by the question itself. We will also focus on cases in which you are only given two options, and forced to choose one. So the complications discussed by Dorst and Mandelkern will not arise.

How do our credences license guesses? Here is a general definition, which puts licensed guesses in terms of rationality:[10]

> **A doxastic state S licenses a guess G so long as it is compatible with maximal rationality that, for the purpose of answering a question {G, G'} truly, an agent in state S guesses G.**

Roughly, you're licensed to guess in whatever way you take to give you the best shot at answering the question truly.[11] (What you guess will be question-dependent: it might make sense to guess G when the alternative is G', but not when the alternative is H.)

I think the definition above is correct. However, sticklers for accuracy-first epistemology — those who want to understand every norm of epistemic rationality in terms of accuracy — might want to get a more complete handle on the notion of accuracy before bringing rationality into the picture. For those purposes we could instead start with the following two stipulations. Assuming that the believers in question have precise credences, the following is true of our

---

[8] Why call it "guessing", then, rather than simply a forced choice under certain constraints? If examples like the ones discussed above bother you, again, you should take my usage of "guess" to be somewhat stipulative. But although this notion of guessing comes apart from our ordinary notion in some cases, I think it is still helpful: a guess is a choice of a *proposition*, with the aim of choosing *truly*. In those respects, the ordinary conception of guessing captures two of the more central ways in which this type of forced choice is distinctive.

[9] Dorst and Mandelkern [2021]. I will come back to their view in chapter 11.

[10] See Builes, Schoenfield, and Horowitz [2022] for a similar definition.

[11] I'll sometimes write questions using the brackets, as I do above, in order to more clearly avoid the complications that arise from imagining them in an actual conversational context.

licensed guesses:

> **Simple Questions:** When faced with a forced choice between two propositions, you are licensed to guess in favor of a proposition in which your credence is highest.

This means that if you are forced to guess between G and G', you are licensed to guess G if your credence in G is at least as high as your credence in G'. Sometimes, you will be licensed to guess either G or G' (when you have equal credence in both).

Simple Questions concerns how we will guess on the basis of our unconditional credences. But we have conditional credences as well, and these can also be more or less accurate. I will also take the following to be true of our licensed guesses:

> **Suppositional Questions:** When faced with a forced choice between two propositions, given some supposition, you are licensed to guess in favor of a proposition in which your conditional credence (conditional on the supposition being true) is highest.[12]

Suppositional questions are just slightly more complicated. Suppose you know that I have a pet, and that my pet is either a cat or a lizard. You know that cats are much more popular to have as pets, and so your credences about what type of pet I have are as follows:

Cr(Cat): 0.8

Cr(Lizard): 0.2

But you also know that lizards are often green, while cats are never green. Consider the following question: supposing my pet is green, is it a cat or a lizard?

Supposing Green: {Cat, Lizard}

Here you will answer on the basis of your conditional credences. Let's suppose those are as follows:

$Cr(Cat|Green) = 0$

$Cr(Lizard|Green) = 0.5$

In response to this suppositional question, your credences license guessing that I have a lizard.

2. A complication: forced choices between two truths or two falsehoods

It's easy to see how this picture works when we consider questions like {Rain, ~Rain}, where the two options form a partition. If Rain is true, ~Rain must be false, and vice versa. Guessing Rain, if it's raining, is obviously more accurate than guessing ~Rain.

It's less easy to see how it works when the options don't form a partition. Suppose it's raining, and it's Tuesday. How should you guess in response to the question {Rain, Tuesday}? Both options are good. And how should you guess in response to the question {~Rain,

---

[12] Some more technical readers might find that the word "supposition" triggers thoughts about imaging. That is not my intention here.

Wednesday}? Both options are bad.

We have already seen some examples where the available options don't form a partition, and where the guesser is pretty sure that both options have the same truth value. Nothing about the definition of a licensed guess depends on the options forming a partition; your choices could fail to be mutually exclusive, fail to be exhaustive, or both. In these cases, just as in cases like {Rain, ~Rain}, your credences license guessing an option just in case you're at least as confident of that option as you are of its alternative. And nothing about the definition of a licensed guess depends on one option being true and the other false.

But here one might raise an objection. How can a true guess count in favor of someone's accuracy when their only option was to guess truly, as with {Rain, Tuesday}? And how can a false guess count against someone's accuracy when their only option was to guess falsely, as with {~Rain, Wednesday}? One way to pose the objection: if both options are true, then you don't deserve any particular credit for guessing truly, since there was no way you could fail. Another way to pose it: if both options are true, you would answer truly no matter what credences you had. How could your response to this question count in favor of *your* accuracy, if it would count in favor of *anyone's* accuracy?

We will answer this question by amending the account of accuracy. While guesses will be licensed in response to any question (including those with two true options or two false options), your answer will only count for or against your accuracy if one of the available options is true and the other is false.[13] This leaves us with a slightly more complicated slogan: **on questions that matter — that is, where one option is true and the other is false — your credences are accurate insofar as they license true guesses. They are inaccurate insofar as they license false guesses.**

This obviously implies that on a question that matters, Anna, who has higher credence in the true option, has more accurate credences than Charlie, who has higher credence in the false option. But what about Beth, who has equal credence in both? Beth will be licensed to guess the true option (which counts in favor of her accuracy) and also licensed to guess the false option (which counts against her accuracy). I propose that we should think of Beth as more accurate (as concerns this question) than Charlie, but less accurate than Anna.

This amendment has a few interesting implications. First, obviously, it means that not every licensed guess will count towards (or against) your accuracy. A guess in favor of some proposition (e.g., "Rain") might count in the context of some questions but not others. And the same question might count in some states of the world but not others. This assumption will come in handy later, especially in Chapter 5, as it will provide a way to argue for certain coherence requirements by focusing on only those states of the world that make a difference for accuracy.[14]

[13] This specification goes beyond the account presented in Horowitz [2019] and Builes, Schoenfield, and Horowitz [2022].

[14] There may be a way to run many of the arguments in this book with slightly different assumptions instead of the ones I make here. For example, in a question between two true propositions, we might count both guesses "equally" towards an agent's accuracy. I suspect that this is a merely verbal difference for my current approach, according to which a guess only counts "for" or "against" one's accuracy in isolation; it may make more of a difference to those who want to extend the guessing framework to create an aggregative scoring rule. Similarly, we might work with different ways of ranking Anna, Beth, and Charlie – for instance, counting Beth just as accurate as Anna or just as accurate as Charlie, rather than between the two. I will leave the full investigation of this question for future research, but note some reasons to prefer my ranking for the time being. First, given a question {P, Q} that matters, Anna is guaranteed to answer the question correctly, while Charlie is guaranteed to answer it incorrectly. Beth might go either way. If you need to choose one of them to go take a quiz containing only question {P, Q}, and you only care about answering truly, it seems clear that you should prefer Anna over Beth, and Beth over Charlie. Second: an

Second, it means that you won't always know which guesses count, because you often won't know whether a given question offers a choice between one true and one false proposition. As an analogy, suppose you are taking a multiple-choice quiz on which only some of the questions will be graded, but you don't know which. If you want to get credit wherever possible, you'll answer each question. When we think about educated guessing, we can imagine a student who is prepared to take such a quiz and wants to get a good grade, but does not know which questions will be on the quiz or which questions will be graded.

You might worry that, if a question only counts when exactly one of the options is true, this could change our guessing behavior. Someone taking a multiple-choice quiz might reason as follows: "I don't know which questions are going to be graded. This question, here, between P and Q, may or may not be — it's possible that P and Q have the same truth value. If they do, my answer won't matter. I might as well answer the question *as if* it's going to count. In other words, I will answer it *on the supposition that exactly one of P or Q is true.*" In order to do well on the quiz, it seems like we should convert any non-suppositional question into a suppositional question to make sure we get "credit" if credit will be awarded. Couldn't this somehow change how we guess?

Fortunately, treating questions as suppositional in this way should plausibly preserve which guesses are licensed. If you're more confident of P than of Q, then you will also be more confident of P than of Q given the supposition that exactly one of those propositions is true. Why is this? Intuitively, the reason is that when you are comparing P and Q, any difference in your confidence must come from possibilities where exactly one of them is true. Your credence in P covers all of those possibilities where P and Q are both true, plus any more possibilities where only P is true; your credence in Q covers all of those possibilities where P and Q are both true, plus any more where only Q is true. Since the (P & Q) possibilities are shared between them, they cannot account for any difference in your credence. If you're more confident of P than of Q, it must be because you're more confident of the only-P possibilities than of the only-Q possibilities. So if you focus only on those possibilities, as we just imagined, you'll still be licensed to guess P.[15]

3. An interesting upshot: the accuracy of conditional credences

I said above that your guesses in response to *suppositional* questions, which we answer on the basis of our conditional credences, could count towards or against our accuracy. Since both conditional and unconditional credences license guesses, this means that the guessing account

---

unpublished project on comparative confidence by Fitelson and McCarthy addresses some of these same questions. (See their [ms] and Fitelson [ms], Chapter 3, for details.) Although their approach differs from mine in a few respects, they agree about how to rank agents like Anna, Beth, and Charlie. Thanks to a reviewer for suggesting these alternative routes.

[15] Here is a more formal proof. We want to know whether it can happen that $Cr(P) > Cr(Q)$, while $Cr(Q|(P \text{ xor } Q)) > Cr(P|(P \text{ xor } Q))$.

$Cr(Q|(P \text{ xor } Q))$ is equivalent to $Cr(P\&{\sim}Q)$. And $Cr(P|(P \text{ xor } Q))$ is equivalent to $Cr(Q\&{\sim}P)$.

But $Cr(P) = Cr(P\&Q) + Cr(P\&{\sim}Q)$, and $Cr(Q) = Cr(Q\&P) + Cr(Q\&{\sim}P)$.

Therefore $Cr(P) > C(Q)$ iff $Cr(P\&\neg Q) > C(Q\&\neg P)$ iff $Cr(P|(P \text{ xor } Q)) > Cr(Q|(P \text{ xor } Q))$.

So it will not happen that you're licensed to guess P in response to {P, Q}, yet licensed to guess Q in response to {P, Q} on the supposition that exactly one of P or Q is true.

This proof relies on some assumptions (probabilism and the ratio formula for conditional credences) that I have not argued for. However, I think the intuitive motivation is strong enough that I am happy to proceed for now on the assumption that changing to conditional credences will not matter. I will leave full discussion of the ratio formula for future research.

allows us to assess the accuracy of conditional credences as well as unconditional credences.

What happens when the condition — or the supposition of a suppositional question — is false? I propose that guesses in response to suppositional questions will only count when the supposition is true. (Moreover, because of the complication above: they will only count in the context of questions where one option is true and the other is false.) So, strictly speaking, the guessing framework only allows us to assess the accuracy of *some* conditional credences. A conditional credence, Cr(H|E), is more accurate insofar as it licenses true guesses when E is true, and less accurate insofar as it licenses false guesses when E is true. If E is false, it is not assessable at all.

Mainstream views in EUT do not tend to assess conditional credences for accuracy. But they do not rule out the possibility, and some ways of assessing conditional credences have been proposed. In my [2021], I defended an approach aligned with the prior suggestion.[16] There I argued that:

For any conditional credence Cr(H|E),

> (1) If ~E, the score of Cr(H|E) is undefined.
>
> (2) If E & H, the score of Cr(H|E) is a function of the distance between Cr(H|E) and 1 (as measured by a qualifying scoring rule).
>
> (3) If E & ~H, the score of Cr(H|E) is a function of the distance between Cr(H|E) and 0 (as measured by a qualifying scoring rule).

I argued for this way of assessing conditional credences by considering a Diachronic Dutch Book. In this Diachronic Dutch Book argument, an agent makes three bets at different times. The first bet, on E, is paid out no matter what. But the second bet, on H, is *conditional* on E, and is called off if E turns out to be false, and pays out if E turns out to be true. The third bet, also on H, is simply never offered if E turns out to be false.

The proposal above is similar, I argued, to a bet that's called off or not offered if E turns out to be false — though in this case, it's the entire accuracy score that's called off or not offered. Intuitively, I argued, there is a difference between a bet that pays out $0 and a bet that is never settled. My suggestion is that we understand conditional credences, where the condition is false, as similar to the latter case.

There are also a few other proposals in the literature for scoring conditional credences. A 2011 blog post by Brian Weatherson suggests a similar approach, applying the Brier score to an agent's conditional probabilities, but only considering those whose conditions are satisfied.[17] And a 1982 paper by Dennis Lindley argues for another similar approach: for Lindley, the score of a conditional credence is just its Brier score when E is true, and 0 otherwise.[18] (He defines it as the Brier score of the credence, multiplied by the truth value of E.) If we are primarily interested in aggregating all of an agent's conditional credences, and looking at the total score, it doesn't matter whether those with false conditions are given a score of 0 or simply undefined. But there is a difference if we look at the conditional credences one by one; I prefer the approach

[16] Horowitz [2021], p. 315. See also Schervish, Seidenfeld, and Kadane [2009] for a similar approach.
[17] Weatherson [2011].
[18] Lindley [1982]. Thanks to Snow Zhang for pointing me to this paper.

suggested in my earlier paper, and separately by Weatherson, for that reason. Lindley's approach treats conditional credences with false conditions like bets that pay out $0, rather than bets that are called off or not offered.

4. Some quick comparisons

My picture of guessing bears some similarities to other, related notions discussed by epistemologists and by philosophers of language and mind. Although I will not attempt to draw out those comparisons in full, I flag a few salient related notions here.

First, guessing is a type of forced choice. Although a "forced choice" is something of a decision theorist's fiction, it approximates real-life choices in which agents are short on time or cognitive resources – or where their options have been constrained by some contextually salient force, as on a multiple-choice test in school. And more relevantly for my purposes, forced-choice cases have widespread use in theorizing: in philosophical contexts as diverse as the ethics of self-defense and the pragmatic encroachment debate, as well as in research in the social and natural sciences.[19] For those who worry that my notion of guessing is too artificial, it may help to think of this project as another application of this commonly-used theoretical tool.

Guessing also has some things in common with betting. In a way, one might see my notion of guessing as a very restricted kind of bet: one where our choices are limited to two-way forced choices between propositions, and the payouts are simply "good" if the guess is true and "bad" if it's false. (This framing does not strike me as inaccurate, though I prefer to think of guessing as a depragmatized notion of betting, rather than a specific type of bet. I discuss this issue further in Chapter 4's discussion of Immodesty.) A perhaps less-obvious difference between betting and guessing is that betting allows us to make very fine-grained and detailed comparisons. Because we can vary both the odds and the payout, that is, betting affords a very rich space of available decisions and preferences. As we will see throughout this book, guessing does not come with the same kind of rich structure. But richness will be required to get certain results. As those come up, I will discuss the possibilities for obtaining richness in other ways, without monetary payouts.

Patrick Maher's *Betting on Theories* puts forth a notion of "acceptance" which has some things in common with guessing – it's an attitude towards a proposition which is in an important sense truth-oriented, distinct from both full belief and high probability. But there are also important differences: Maher associates acceptance with "sincere intentional assertion", rather than forced choice. Maher also admits a wider range of considerations which would contribute to "cognitive utility", determining whether acceptance is rational; these include logical strength (or "informativeness") and verisimilitude.[20]

Another similar idea is "weak belief", defended by Hawthorne, Rothschild, and Specter, as well as a related account of guessing that has been discussed in recent papers by Ben Holguín, Kevin Dorst, Matt Mandelkern, and others.[21] (I discuss the connections here in a bit more detail in Chapter 11.) Like my notion of guessing, the attitudes these authors are interested in are

[19] For self-defense in particular, see, e.g., Montague [1981]. For pragmatic encroachment, see Schroeder [2017]. Forced-choice cases are of course widespread in applied ethics: just think about trolley problems. It would be silly to try to catalog examples of the uses of forced-choice tasks in the sciences, but a quick Google scholar search will reveal thousands of results.

[20] See Maher [1993], especially Chapter 6.

[21] See, e.g., Hawthorne, Rothschild, and Spector [2016] for weak belief. For recent accounts of guessing, see Holguín [2021], Dorst and Mandelkern [2021], Dorst and Mandelkern [2022], and Holguín and Goodman [2022].

question-relative. However, these authors do not restrict attention to two-option cases, and like Maher, they do not presuppose that the guesser's only aim is to guess truly. And unlike many of these authors, I am not committed to the view that guessing is believing.

Easwaran [2016]'s account of full belief is also notable. Easwaran is motivated by one of the same things that motivate my project: providing a straightforward way in which credences can be said to "aim at truth". But his project also uses epistemic utility theory, and like the accounts mentioned above, is committed to controversial claims about full belief, about which I would like to stay neutral.

Finally, I should mention the work on comparative confidence by Branden Fitelson, together with David McCarthy and Kenny Easwaran.[22] Fitelson builds an accuracy-based account of coherence norms for comparative confidence. The view he develops has some marked differences from the one I discuss here – for one, it uses a quantitative notion of accuracy, in line with epistemic utility theory – but many commonalities as well. In what follows, I note places where my work has especially benefited from Fitelson's (especially in Chapter 5).

5. How we will test this account

This chapter has set out the basic building blocks for the guessing account of accuracy. To recap:

> **The main idea:** Your credences are accurate insofar as they license true (instead of false) guesses, and inaccurate insofar as they license false (instead of true) guesses. Suppositional questions (and by extension, conditional credences) will only count towards or against your accuracy when the supposition is true.

The rest of the book tests and extends this idea. In the next chapter I will argue that this account can meet what I see as two basic requirements on any acceptable notion of accuracy: (a) immodesty, or the sense in which a rational agent must see her own credences as her best option, accuracy-wise; and (b) the truism that credences are more accurate as they get closer to the truth. The following several chapters show how we can use the guessing account to vindicate other formal norms. The very end of the book will discuss connections between guessing and other important notions in epistemology, as well as the philosophy of mind, language, and action.

[22] Fitelson [2016] *Coherence*, ms. Available at https://fitelson.org/coherence/coherence_duke.pdf. See also Fitelson and McCarthy [2014], ms.

Chapter 4. Truth-Directedness and Immodesty

Before we put this account of accuracy to work, we need to make sure it meets two conditions. I see the following as two theses that a view of accuracy *must* be able to make sense of, in order to count as an account of accuracy (rather than something else) at all. These are, roughly:

> **Truth-Directedness.** It is better, accuracy-wise, to have one's credence closer to 1 in truths, and closer to 0 in falsehoods.
>
> **Immodesty.** A rational agent must regard her own credences as best, prospectively, in terms of accuracy.

The first condition sets out a connection between credences and truth; the second one is a connection between *rational* credences and truth. (I'll say more in what follows, both about what these principles mean and why we should want to support them.)

In my [2019], I gave an explicit argument for Immodesty, and a brief and sketchy argument for Truth-Directedness. This chapter both improves on and complicates these arguments, drawing on subsequent work by Gabrielle Kerbel, and conversation with Mike Titelbaum. To forecast, I will argue that we can support a plausible version of Truth-Directedness. We can also support Immodesty, but only with the help of an assumption about the richness of our credal domain.

1. Truth-directedness

The first condition that a theory of accuracy must meet, if it is to count as a theory of accuracy, is that credences *closer to the truth* — closer to 1 if a proposition is true, and closer to 0 if it is false — must count as more accurate. Following Joyce, let's call this general thought "truth-directedness." Epistemic utility theory clearly delivers truth-directedness: on that view, it is axiomatic that accuracy is a function of the distance between a credence and a truth value, which gets better as the distance decreases and worse as it increases. How can the guessing account meet this desideratum?

I will argue here that it can — but in a limited way. The guessing account does not vindicate the strongest version of truth-directedness, which some iterations of epistemic utility theory (such as Joyce [1998]) build in as an axiom.[1] But it does vindicate some nearby principles. In this subsection I will build up the guessing account's view of the relationship between "closeness" and accuracy, and then explain why it ends up looking different from that of epistemic utility theory.[2]

---

[1] On other iterations of epistemic utility theory, strict propriety – a version of Immodesty – is axiomatic instead, and truth-directedness follows from strict propriety. Thanks to a reviewer for helpful clarification on the dialectic here.

[2] To head off a possible worry: vindicating Truth-Directedness does not involve taking any sort of stand on what other properties might be valuable in a belief state. For instance, one might think that beyond mere truth, it is good to have knowledge, understanding, reliability, or rationality. As I see it, Truth-Directedness is compatible with any position in those other debates about epistemic value. (In many cases, certain positions in those debates amount to the view that something *other than accuracy* is valuable. You might think that in some cases it's *better* to be rational than to be accurate: for example, when you have misleading evidence. But Truth-Directedness is only a thesis about accuracy, not about *overall* "bettterness".)

1.1 The main idea: look at many different questions.

Joyce [1998] expresses skepticism that guesses could deliver Truth-Directedness. Following Jeffrey [1986], he writes, "[w]hen one tries to guess, say, the number of hits that a baseball player will get in his next ten at-bats, one aims to get the value exactly right. Guessing two hits when the batter gets three is just as wrong as guessing two hits when he gets ten. In guessing, closeness does not count."[3] By contrast, *estimation* can be more nuanced and graded. In that vein, Jeffrey writes: "[e]stimation is not a matter of trying to guess the truth value, e.g., 2.4 might be an eminently reasonable estimate of the number of someone's children, but it would be a ridiculous guess."[4]

While Joyce and Jeffrey's remarks about guessing seem right, they do not threaten the project here. As I understand guesses, we never simply guess about one proposition: it is always in contrast to an alternative, and I assume that agents have credences as well as guesses. (If Joyce and Jeffrey's main aim is to argue that we ought to have credences *rather than* simply having guesses, what I am saying here is compatible with their view.) On my view, it's implausible that we would be licensed to guess that someone has 2.4 children — unless the only alternative presented to us is some other obviously false proposition. Instead, someone who *estimates* that a person has 2.4 children might guess that the person has two when offered the choice {two children, four children}, for example. And here, closeness would count, insofar as one's credences licensed a true guess rather than a false one.

But even when we consider comparisons, it might still seem that the guessing account can't possibly support Truth-Directedness. Suppose it's raining, and you and I are both asked to choose between {Rain, ~Rain}. We both have greater credence in Rain than in ~Rain, but your credence in Rain is closer to 1: it's 0.8, say, and mine is 0.6. It seems like your credence in rain is more accurate than mine. But in response to the question {Rain, ~Rain}, we will make the same guess. So how can our accuracy differ?

What's important to notice here is that a person's credences will license *many different* guesses in response to different questions. And although in each case the accuracy of a single guess is all-or-nothing — it's true that a 0.8 credence in rain will do just as well as a 0.6 credence in rain, *when the question is {Rain, ~Rain}* — when we look at many questions, we can begin to see more differences. This gives us the beginning of an argument for Truth-Directedness. If it's raining, your credence is more accurate than mine *not* in virtue of our guesses in response to {Rain, ~Rain}, but in virtue of our guesses in response to other questions. We can assess the accuracy of a given degree of credence by seeing what guesses it licenses across many different contexts.

Interestingly, a similar issue arises for EUT. Joyce [1998] imagines an objector pointing out that his notion of gradational accuracy, which depends only on an agent's credences in various propositions and the truth values of those propositions, is unable to capture "verisimilitude". Specifically, it cannot explain the sense in which Kepler, who believed the Earth's orbit is elliptical, had more accurate beliefs than Copernicus, who believed it was circular. (Both astronomers' theories are false, but since the Earth's orbit is *nearly* elliptical, Kepler's theory was closer to the truth.) Joyce agrees with the objector that looking at this one attitude — supposing each astronomer had the same credence in his preferred theory — cannot bear out the thought that Kepler was more accurate. But if we look at their credences in related propositions, we will likely see that Kepler does better. For example, Kepler, but not Copernicus,

[3] Joyce [1998], p. 587. See also Joyce [2009], p. 269, fn. 8.

[4] Jeffrey [1986], p. 51.

would have high confidence that the average distance between the earth and the sun is different in different seasons. We can still make sense of the thought that Kepler's theory is more accurate, even though the difference does not show up when we look directly at their credences in their respective theories.[5]

Our task now is to show that a 0.8 credence in Rain is more accurate than a 0.6 credence in Rain by looking at many different questions, and the guesses that each credence licenses in different cases. I'll look at two ways of spelling this strategy out: a personal approach, which looks at a particular agent's credences, and an impersonal approach, which views credences as abstract mathematical objects.

1.2 Truth-Directedness: the personal approach

How might we show that for some particular agent, it's more accurate to have 0.8 credence in Rain rather than 0.6? In my [2019] I gave a rough argument. I imagined a continuum of possible coins, weighted towards Heads to different degrees, and proposed a question that pitted your guess about Rain against an arbitrary weighted coin's "guess". I argued that you should expect to "beat" a given coin if you are more opinionated than the coin. So as your credence in a true proposition gets higher, you should expect to beat more and more possible coins.

Gabrielle Kerbel [ms a] proposes a clever and more precise way to spell out this thought. Her argument is, roughly, as follows. Consider the following scenario:

> **Coin Game**: Matthew and Susan are playing a guessing game. Each gets a point if they guess correctly. Today, they are answering the question {Heads, P}. P is true. Susan's credence in P is 0.8, Matthew's credence in P is 0.6, and the coin in question is weighted 0.7:0.3 towards Heads.

Intuitively, and according to Truth-Directedness, Susan's credence in P is more accurate than Matthew's. And if both Susan and Matthew obey the Principal Principle, matching their credences to known chances, we expect this difference to come out in a case like **Coin Game**. Since Susan and Matthew both know the coin's weighting, they will both assign 0.7 to Heads. In response to {Heads, P}, then, Susan will guess "P" and Matthew will guess "Heads".

So far we have shown that Susan and Matthew will guess differently on this question, but not yet that Susan's credence is more accurate. In fact, the most likely outcome in the story above is that the coin will come up Heads, and both Susan and Matthew's guesses will be true. However, in the long run, we should *expect* that the coin will sometimes come up Tails. Therefore, we should expect that there are some cases in which Susan's guesses *are* better than Matthews — and none in which Matthew's are better than Susan's. (Since are assuming that P is true, there is no possible scenario in which Matthew guesses truly, but Susan guesses falsely, in response to {Heads, P}.) Kerbel proves that in general, if we consider different possible coins with arbitrary weightings, we will expect Susan to outperform Matthew in this game: she will do better at guessing truly, in expectation. Kerbel goes on to show that we can extend this thought to create a strictly proper scoring rule for guessing.[6]

---

[5] Joyce [1998], p. 592. Strictly speaking, the situation here is not exactly analogous with the one that confronts the guessing framework. Joyce is explaining the superior accuracy of Kepler's *entire belief state*, by looking at credences in different propositions. By contrast, I am looking at how *one* credence performs across many different circumstances.

[6] This is a very interesting result that undoubtedly deserves more space than I give it here. However, as of this writing, Kerbel's work is preliminary and subject to change; furthermore, it takes the kind of aggregative approach

A nice feature of Kerbel's argument is that the expectation is not relative to a person's own idiosyncratic credences, but appeals instead to the probability of guessing P or Heads given the setup of the exam. Her conclusion could also be stated like this: on average, you will do better on this exam as your credence in truths gets higher, and as your credence in falsehoods gets lower.

But in another respect, Kerbel's argument *is* specific to an agent's credences. Because the argument requires that Susan and Matthew each guess regarding the question {Heads, P}, it only works if Susan and Matthew's credences in Heads obey the setup of the case. In order to make sure Susan and Matthew both treat information about the coin in the intended way, we must assume that they obey the Principal Principle, and both must countenance the possibility that any possible coin could pop up on the exam. In the next subsection I'll look at how the argument runs into problems when those assumptions are relaxed.

1.3 A complication for the personal approach: irrational agents, gappy credences

The problem with the personal approach, as spelled out above, is that it does not apply to everyone. Agents who do not follow the Principal Principle are one sort of exception, and agents with "gappy" or incomplete credence functions are another.

Let's consider gappy credences first. Imagine two very simple believers, Tim and Jim. Tim and Jim each only have credences about whether it's raining or not. Tim has 0.8 credence in Rain, and 0.2 credence in ~Rain, and Jim has 0.6 credence in Rain and 0.4 credence in ~Rain. If it is in fact raining, are Tim's credences more accurate than Jim's? It seems that intuitively, we should say "yes", but the guessing approach seems to say "no". Tim and Jim's credences both license exactly one guess: when given the question {Rain, ~Rain}, they both license guessing Rain. Tim and Jim have different credences, but no difference in guesses — which means, according to the guessing framework, no difference in accuracy.

A similar problem arises, in a more complicated form, when we consider agents who do not follow the Principal Principle.[7] Let's return to **Coin Game**, from the previous subsection, where we expected Susan to do better than Matthew regarding the question {P, Heads}. We saw that Susan and Matthew's guesses will only make a difference to their accuracy on those occasions where Heads is actually false. To sum up, the difference-making cases are those with the following features:

(a) $Cr_{Susan}(P) > Cr_{Susan}(Heads)$ *so, Susan will guess P*;

(b) $Cr_{Matthew}(P) < Cr_{Matthew}(Heads)$ *so, Matthew will guess Heads*; and

(c) P is true, but Heads is false.

As the argument goes, Susan is more accurate than Matthew because of cases with features (a) - (c). But why are we so sure that such cases will ever come about? If we are allowed to flip the coin as many times as we want, (c) will probably happen; we can argue, at least, that we expect it will happen. On the other hand, (a) and (b) only seemed plausible because we assumed that both

---

that I am trying to avoid if possible. Readers who are interested in combining the approach I take here with the decision-theoretic strategies of epistemic utility theory are encouraged to look up her work.

[7] Thanks to Mike Titelbaum for raising the point discussed in this section.

Susan and Matthew would match their credence in Heads to the known chance of heads.

Consider another version of Matthew — call him Lucky Matthew — who does not follow the Principal Principle. Let's grant that this is irrational. But although he's irrational, Lucky Matthew is lucky: in particular, he is mysteriously accurate regarding Heads. Although he knows that the chance of heads is 0.7, he does not set his credence in Heads to 0.7. Instead, whenever the coin will *in fact* land heads, Lucky Matthew adopts credence 0.7 in Heads (and keeps his 0.6 credence in P). Whenever the coin will land tails, Matthew adopts credence 0.5 in Heads (and keeps his 0.6 credence in P). (This isn't because Lucky Matthew has some "inadmissible evidence" that allows him to screen off his information about chance. He's just very lucky.) Because he is very lucky in this way, Lucky Matthew will *always* answer the question {Heads, P} correctly. And this means that he is just as accurate as Susan. The upshot of this is that, although it still seems that Lucky Matthew's 0.6 credence in P is less accurate than Susan's 0.8 credence in P, we can't use **Coin Game** show that it is.

"Fine," one might respond, "let's just find some more coins. Matthew is bound to guess wrong at some point." But this doesn't work. Imagine an even luckier version of Matthew, Very Lucky Matthew. Very Lucky Matthew isn't omniscient; his credences in truths are often less than 1, and his credences in falsehoods are often more than 0. But Very Lucky Matthew is mysteriously reliable about *everything*, in the following sense: his credence in any true proposition is always (luckily!) higher than his credence in any false proposition. Now compare Very Lucky Matthew to Omniscient Susan, who has credence 1 in all truths, and credence 0 in all falsehoods. Truth-Directedness says that Omniscient Susan is more accurate than Very Lucky Matthew. But it seems that we can't do justice to this verdict within the guessing framework: Very Lucky Matthew and Omniscient Susan will never differ in the content, and hence the accuracy, of their guesses. Every time Omniscient Susan guesses truly, Very Lucky Matthew will guess truly as well. And every time Very Lucky Matthew guesses falsely (that is, in response to questions with two false options), so will Omniscient Susan.

What this objection shows is that, as far as the personal approach goes, the guessing account cannot support the strongest version of Truth-Directedness. That is, on the guessing account, although a weaker principle is true:

> **Comparative Truth-Directedness:** For any true proposition P and false proposition Q, a person's credences in P and Q will always be more accurate when $Cr(P) > Cr(Q)$ than when $Cr(P) \leq Cr(Q)$.

the following, stronger principle appears to be false:

> **Non-Comparative Truth-Directedness:** For any true proposition P, a person's credence in P becomes more accurate as it gets closer to 1 (and for any false proposition P, a person's credence in P becomes more accurate as it gets closer to 0).

Although Kerbel's argument can support Non-Comparative Truth-Directedness for agents who also obey the Principal Principle, without the Principal Principle we can only support Comparative Truth-Directedness.

In fact, Comparative Truth-Directedness was argued for in the last chapter: there, I said that when we consider questions which contribute to or detract from one's accuracy — those with one true option and one false option — one's credences are more accurate if they *only* license guessing in favor of the truth, and less accurate if they license guessing in favor of the falsehood.

Kerbel's response to this problem is to retreat to a different weakening of Truth-Directedness. She argues for the following:

> **Truth-Directed Safety and Promise:** For question {P, Q} with true proposition Q, if Cr(Q) is closer to the truth than Cr'(Q) at world *w*, then upon receiving evidence that bears on P, but not Q, Cr licenses both *safer* and *more promising* guesses at *w* than Cr'.[8]

The idea here is that when you receive evidence relevant to P (the false proposition), you should update by conditionalization. When you change your credence in P, it might change your guess about {P, Q}. No matter whether your evidence about P is misleading or non-misleading, holding Cr(Q) will leave you in at least as good a position, and sometimes better, than Cr'(Q), with respect to the question {P, Q}.

First take the case where the evidence is *misleading*, and so updating on that evidence will lead to your increasing credence in P. And suppose your initial credence in Q is higher than your credence in P, so you start off with a true guess about {P, Q}. In this case, it will take more misleading evidence to flip your true guess to a false one if you hold Cr than if you hold Cr'. (And if your initial credence in Q is not higher than your initial credence in P, you will be no worse off holding Cr than Cr'.) This is the sense in which Cr is *safer* than Cr'.

Second, suppose the evidence is *non-misleading*, so updating will lead you to lower your credence in P. And suppose your initial credence in Q is *lower* than your initial credence in P, so you start off with a false guess about {P, Q}. In this case, it will take less non-misleading evidence to flip your false guess to a true one if you hold Cr than if you hold Cr'. (If your initial credence in Q is not lower than your initial credence in P, you will be no worse off holding Cr than Cr'.) This is the sense in which Cr is *more promising* than Cr'.

For an individual considering the accuracy of her own credences, Safety and Promise do seem to get us some of what we wanted out of Truth-Directedness. Suppose the accuracy fairy comes down from epistemic heaven and offers you a favor: she will swap out some of your credences to make them closer to the truth values in the relevant propositions. Epistemic utility theory says that if your only concern is accuracy, you *should* accept the favor: you'll *definitely* become more accurate. The guessing framework can't say that, but it can say that the fairy's deal is a good one: it can only help, and can't hurt.

However, this doesn't get us everything we wanted. We still can't say that Tim is more accurate than Jim, supposing we know that their credences will remain sparse and gappy (or morbidly, supposing that Tim and Jim are on their deathbeds, and their time to accommodate new evidence has come to an end). And we still know that Very Lucky Matthew has no reason to take the accuracy fairy's offer, unless he's worried that his luck will run out. These problems motivate taking a different approach to Truth-Directedness.

1.4 Truth-Directedness: The impersonal approach

This second approach, which I will call the "impersonal approach", maintains the same central strategy for supporting Truth-Directedness: looking at a credence in P in comparison to other credences, rather than in isolation. But it gives up the idea that we can use an arbitrary person's credences to frame our argument. Instead, we will look at credences as abstract objects, separate from any one person's belief state, and look at how those credences fare in different

[8] Kerbel [ms c], p. 23-24; Kerbel discusses Safety and Promise separately.

environments.

If we decouple a level of credence from the rest of the credence function, we can support the following weak dominance principle:[9]

> **Weak Truth-Directed Dominance:** For any true proposition P, and two credences in P a and b such that a > b, a is more accurate than b in the following sense:
>
> For any possible credence n in a false proposition, Q, the combination of a and n licenses at least as many true guesses as the combination of b and n, and for some values of n, the combination of a and n licenses strictly more true guesses than the combination of b and n.

To see why this is true, suppose $a = 0.8$ and $b = 0.6$. Now consider the possible values n could take, and pair those with a and with b.

(a) $n > 0.8$: both combinations will only license guessing falsely.

(b) $n = 0.8$: the *a*-combination will license guessing truly or falsely, while the *b*-combination will license only guessing falsely.

(c) $0.6 < n < 0.8$: the *a*-combination will license only guessing truly, while the *b*-combination will license only guessing falsely.

(d) $n = 0.6$: the *a*-combination will license only guessing truly, while the *b*-combination will license guessing truly or falsely.

(e) $0.6 < n$: both combinations will license guessing falsely.

In each case, whenever the *b*-combination licenses a true guess, so does the *a*-combination. But in some cases, the *a*-combination only licenses a true guess while the *b*-combination only licenses a false guess.

How is this different from the argument above, in **Coin Game**? That strategy attempted to find a situation like (c) for Susan and Matthew, who had 0.8 and 0.6 credence in P, respectively. But while we might *expect* situations like (c) to crop up for most pairs of believers, given certain other plausible assumptions, such situations are not inevitable. As we saw, a person might have their credences arranged like Very Lucky Matthew, or they might have gaps in their credences like Tim and Jim.

If we depersonalize the argument, however, we don't need to rely on finding situations like (c) in the wild. We can instead look at them as abstract mathematical possibilities.

Is this enough? It's certainly weaker than the form of Truth-Directedness that appears in EUT. For example, Joyce [2009] takes the following to be axiomatic:[10]

> **Truth-Directedness:** For credence functions b and b*, if b's truth-value estimates are

---

[9] Thanks to David Christensen for suggesting this way of going. This approach is consistent with Kerbel's, but does not appeal to future changes in response to evidence.

[10] Joyce [2009] p. 269.

uniformly closer to the truth than b*'s at world ν, so that either $b_n* \geq b_n \geq \nu_n$ or $b_n* \leq b_n \leq \nu_n$ for all n and $b_n* > b_n \geq \nu_n$ or $b_n* < b_n \leq \nu_n$ for some n, then $S(b*, \nu) > S(b, \nu)$.

Note that Joyce's version is both non-comparative and "personal", in the sense that it deals with an entire credence function rather than individual values for propositions. The guessing account can't establish this. However, this should not come as too much of a surprise. This version of Truth-Directedness is the starting point for EUT. Since the guessing framework starts with comparative judgments, rather than a gradational notion of accuracy, it makes sense that it would end up with something weaker.

Is it a problem that we can't immediately say that Tim, holding his gappy credences fixed, is more accurate than Jim? Maybe not. Suppose I argued that in fact, \$2 isn't worth any more than \$1, because there is only one bottle of Coke in the vending machine. This might be right in some limited sense — maybe *I'm* no better off having \$2 than I would be having \$1, in this situation — but it's obviously false in a wider sense. Similarly, Tim doesn't benefit from having a higher credence in Rain than Jim does, because his credences are gappy. But it's still true that a 0.8 credence in Rain is more accurate than 0.6.

The depersonalized approach is also in keeping with the idealizations we need to make to get the guessing framework up and running in the first place. We already have to say that Tim and Jim's accuracy depends on their guesses, *regardless of whether they actually make any guesses*. And we are already committed to only counting those guesses that are rationalized by the aim of guessing truly, *regardless of whether the agent ever has any such aim*. In this way, the guessing approach already abstracts away from a particular person's circumstances and psychology. Considering a certain credence value outside the context of any one person's belief state could be seen as another idealization of the same type.[11]

2. Immodesty

Beliefs and truth have a special relationship: to believe some proposition is to regard it as true. If credences play a similar role to beliefs, then they should have a similar relationship to accuracy. Having a credence amounts to taking a stand on accuracy just as having a belief amounts to taking a stand on truth. For that reason, rational credences should be *immodest*. That is, they should regard themselves as having the best available prospects for accuracy. If we find that some credence function regards itself as less accurate than some alternative (specific, rigidly designated) credence function — or if we find that it regards itself and some alternative credence function as equally accurate — that's a sign that either there's something wrong with our credences, or something wrong with our way of assessing accuracy.

Some people might find this way of framing Immodesty to be contentious. They might argue that rational credences should see themselves as best in *some sense*, but that whether that sense is best understood as *accuracy* is up for debate. For example, if one thinks that there is more to rational belief than just accuracy — clutter avoidance, fruitfulness, or some other virtue — one might think that immodesty is true, but is best spelled out in terms of accuracy *and* that other virtue. Alternatively, one might think that Immodesty is true, but that it tracks some other feature entirely. Gibbard [2008], for example, argues Immodesty tracks "guidance value".

I find the purely accuracy-centered version of Immodesty to be the most compelling, and I adopt it as a starting point here. In fact, I find it plausible that (accuracy-based) Immodesty is a

[11] Thanks to David Christensen for suggesting this line of argument; see his [1996] for a similar point, made in the context of the Dutch Book argument. See Horowitz [2024] for further discussion of the other idealizations.

constitutive requirement of beliefs and credences, such that a mental state does not qualify as a belief or credence if it does not take itself to be uniquely prospectively best in terms of accuracy. (This is stronger than the idea that Immodesty is required for *rational* credence.) A consequence of this stronger view is that if the guessing account can support Immodesty, then guessing can help explain what is distinctive about doxastic states, as opposed to other mental states.

2.1 The general case for Immodesty

The name "immodesty" comes from David Lewis's 1971 paper, "Immodest Inductive Methods." In his terminology, an immodest inductive method is one that estimates its own accuracy as best, compared to other inductive methods. Lewis offers the following as an analogy, which both illustrates what Immodesty is and why trustworthy methods must be immodest. Suppose Consumer Reports, Consumer Bulletin, and other product-reviewing magazines were to review product reviewing magazines. If a magazine is to be trusted, Lewis argues, it must at the very least recommend itself.

Why is that? Well, suppose *Consumer Reports* recommended *Consumer Bulletin*. (In this case *Consumer Reports* would be "modest" in Lewis's terms.) Since the two magazines sometimes disagree about which products are best (vacuum cleaners, toasters, etc.), *Consumer Reports* would thereby be giving inconsistent recommendations. On p. 3, say, *Consumer Reports* recommends the Toasty Plus as the best toaster, specifically noting that it is better than the Crispy Supreme. On p. 7, in its reviews of product reviewing magazines, it recommends *Consumer Bulletin*. Then, when you open up *Consumer Bulletin* to the toaster reviews, you find out that *Consumer Bulletin* recommends the Crispy Supreme, specifically noting that it is better than the Toasty Plus. *Consumer Reports* can't be trusted if its advice is incoherent in this way.

Of course, this does not mean that any magazine which *does* recommend itself is trustworthy. Self-recommendation is a necessary, but not a sufficient, condition for trustworthiness. Furthermore, self-recommendation is only necessary when the alternatives are similar in certain respects: in the case of product reviewing magazines, the magazines must rate the same products and have access to the same information. If the *Consumer Reports* team knew that another magazine with a bigger budget was able to conduct better tests, or rate more products, *Consumer Reports* might be able to recommend that magazine without incoherence.

Unlike consumer magazines, our credences do not make recommendations. But our accuracy measures do: they deliver a judgment about how accurate different belief states are in different states of the world. And our credences, in combination with our accuracy measures, give a particular kind of recommendation: they tell us which doxastic state is best, accuracy-wise, from our own point of view. Carrying over the lesson about *Consumer Reports*, we can think of that combination of credences and accuracy measures as our own rating system, which we can use to rate various belief states. Immodesty is the thesis that, if we have rational credences and an acceptable way of understanding accuracy, *our own* credences should come out best when assessed in this way. (Just as before, the requirement only applies when we hold fixed certain things about the target belief states: they must range over the same propositions, for example. And they must be specified "numerically", rather than by description. There is nothing wrong with judging that the beliefs of an omniscient God — specified only under that description — would be more accurate than one's own.)

Upon hearing this argument, people often ask whether immodesty requires that we merely rate our own beliefs as *among* the best, or whether we must rate them as *uniquely* best. We can call these options "weak" and "strict" immodesty, respectively. The question arises because the *Consumer Reports* argument is really an argument *against modesty*, which leaves both options open. It's easy to grant the conclusion that I should not deem another belief state *more accurate* than mine. But what's wrong with regarding another belief state and my own as *equally accurate*?[12]

I will argue for strict immodesty. I think there is ample intuitive motivation for this. In addition, Catrin Campbell-Moore and Ben Levinstein have proved that in the EUT framework, once you have "weak immodesty" plus Truth-Directedness, you can derive "strict immodesty" (which delivers the "uniquely best" verdict).[13]

The intuitive case is as follows. Suppose I believe that it's raining. To argue against modesty, we can imagine that my way of assessing accuracy deems a belief in ~Rain *more accurate* than a belief in Rain, given my evidence. This is intuitively an irrational state to be in: if I think ~Rain is a more accurate belief, why not go ahead and adopt it? A person in such a state must either have irrational beliefs or a defective conception of accuracy. And though it is a little less straightforward, we can extend this argument to tell against weak immodesty as well. Suppose I believe it's raining, and my way of assessing accuracy deems a belief in ~Rain *just as accurate* as a belief in Rain. In this case it doesn't seem like there is pressure for me to adopt the belief in ~Rain; doing so wouldn't be an improvement. But there is also no pressure *against* adopting that belief. If the two are equally accurate, why not pick whichever belief I like better on other grounds, or flip-flop between them if I feel like it? This line of thought gives us a first argument against weak immodesty: weak immodesty makes our beliefs look arbitrary, and seems to permit arbitrary changes of belief.

For someone who is weakly immodest in the way described above, it also seems intuitive that really, the thing to do is to suspend judgment. If you see the belief in Rain and the belief in ~Rain as equally accurate (or having an equal shot at accuracy), you should not adopt either of those beliefs. Going back to Lewis's way of framing the issue, we can think of a rational agent as someone who (among other things) adopts the belief state that her own belief state, in combination with her accuracy measure, "recommends". If belief in Rain and belief in ~Rain come with equally high recommendation, suspension of judgment should be recommended more strongly than either.

This argument crucially relies on the thought that we *can* suspend judgment; our options are not limited to belief and disbelief. In a degreed-belief setting, arguing for strict immodesty relies on the thought all sorts of intermediate levels of credence could be available to you: if you deem credences a and c equally good in terms of accuracy, there must some intermediate credence b that you could adopt instead, which should look even better accuracy-wise.

In this way, beliefs and credences are very different from actions. In the practical realm, sometimes our options are limited: even if we are unsure whether it's raining, there is no halfway point between bringing the umbrella or leaving it at home. So it can often be perfectly rational to take an action arbitrarily, while viewing some other action as equally choiceworthy.[14] Weak immodesty is a reasonable constraint on practical rationality. But since our doxastic state can admit more shades of gray, credences should be strictly immodest.

I think of Immodesty as a constraint on acceptable accounts of accuracy. Many epistemic utility theorists do, too: they often build "strict propriety", or the requirement that rational

[12] In conversation, many people have suggested to me that they find non-strict Immodesty to be an attractive position. For my own part, I just can't see the appeal. But see, for example, Daoust [2020] for a defense.
[13] Campbell-Moore and Levinstein [2021].
[14] I am setting aside worries about Buridan's ass; I assume that it is possible to act arbitrarily.

credences be uniquely self-recommending, into their requirements for measures of accuracy. (Epistemic utility theorists also often build in "convexity", which implies that if two credences are equally accurate, there must be another between them that is more accurate still. This thought is similar to the argument against weak immodesty given above.) The motivating idea is that accuracy needs to be able to play this role; if it's what we aim for in our beliefs, it had better turn out that rational beliefs see themselves as fulfilling this aim the best they can.

But while immodesty is primarily a constraint on accuracy measures, it goes hand in hand with our theory of rationality. *Rational* credences are the ones which (together with a good account of accuracy) must self-recommend. Irrational credences might not. In epistemic utility theory, the relevant rational constraint is probabilism, or the requirement that rational credences be probabilistically coherent. In my argument for Immodesty, I'll first rely on the assumption that rational credences are sufficiently "rich". One way to get this result is to assume that rational agents follow the Principal Principle, with some other specifications. Echoing the discussion of Truth-Directedness, I will call this the "personal" argument for Immodesty, since it defends Immodesty within the context of a particular fully-specified credence function. I will then consider the possibility of giving an "impersonal" argument for Immodesty. Although that approach worked in the case of Truth-Directedness, it does not work here; still, it is instructive to see why.

2.2 Immodesty: the personal approach

First, let's get clear on what we are arguing for. We want to show that rational credences are uniquely self-recommending, accuracy-wise, using the guessing account of accuracy. And in particular, we want to show that this is true of a *rational agent's* credences. So we want to show that the following principle is true:

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

The argument I will give is inspired by Allan Gibbard's [2008] argument for immodesty, which is given in practical terms. Drawing on a proof by Mark Schervish, Gibbard imagines a rational agent preparing for a hypothetical series of bets on different propositions, at varying odds. She can choose to bet according to her own expected utilities (those corresponding to her current credences) or others (corresponding to other credences). In order to make good bets, thereby maximizing her expected winnings, she should use her own credences. This shows that she is immodest, in a certain way: she sees her own credences as best for the purposes of making practical decisions.[15]

The guessing argument I will give is a depragmatized version of Gibbard's, in the spirit of (though perhaps not quite as depragmatized as) Christensen [1996]'s depragmatized Dutch Book argument for probabilism. Instead of looking at bets, which pay out in money or utility, we will look at guesses in response to questions. While Gibbard asks us to imagine ourselves choosing credences to use for betting, we will imagine ourselves choosing credences to use for guessing. If Immodesty is true, this means that we will take our own credences to license the best

[15] In giving this argument, Gibbard rejects the possibility of a purely alethic, or truth-directed, version of immodesty. His foil is epistemic utility theory. He does not think that the utility functions (or "scoring rules") used by EUT are best thought of as measures of accuracy; in order to produce plausible results, they must follow axioms that are, in his view, aiming at something other than just the truth. This is why he looks to practical value instead.

guesses.

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 comparing two propositions (you don't know which ones, 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, and you will not have the option of refusing to answer. (In other words, you'll give your guess, in response to a forced choice.) 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 answer using the guesses licensed by that credence function. Which one should you choose? Which one looks prospectively best to you, from your current perspective? For concreteness, let's call your *current* credence function "$Cr_{Current}$", and the credence function that *looks best to you*, prospectively, "Cr_{Best}". (Cr_{Best} is the credence function that $Cr_{Current}$ "recommends".) Immodesty is the claim that $Cr_{Current}$ uniquely recommends itself: $Cr_{Current} = Cr_{Best}$. You should pick your own credences as the best credences to use for guessing.

It is easy to see that your current credences constrain your choices for Cr_{Best}. They have opinions about which guesses are best, accuracy-wise (in response to various forced choices), and so they will want to pick Cr_{Best} that delivers those answers. For example, let's look back at Question 1, from the last chapter:

Question 1: Is it raining, or not raining?

Suppose $Cr_{Current}$ is more confident of Rain than of ~Rain. So in response to Question 1, $Cr_{Current}$ licenses guessing Rain, and does not license guessing ~Rain. $Cr_{Current}$, in conjunction with the guessing account of accuracy, therefore *recommends* guessing Rain in response to this question. By extension, it recommends having higher confidence in Rain than in ~Rain.

Immediately we can see that whatever credence function $Cr_{Current}$ recommends, it must make the same comparative judgments as $Cr_{Current}$ itself. Whenever $Cr_{Current}$ is more confident of one proposition than another, Cr_{Best} must be as well. We can sum this up in the following principle:

> **Comparative Immodesty**: A rational agent should take credences which agree with her own in terms of all of their comparative judgments to be best, by her own current lights, for the purposes of *making true guesses*.

But this does not yet show that $Cr_{Current} = Cr_{Best}$. In order to get the right answer to questions like {Rain, ~Rain}, Cr_{Best} could take any number of forms; for instance, it could just assign 1 to Rain and 0 to ~Rain. And more generally, guaranteeing that two credences agree on their comparative judgments does not guarantee that they will agree numerically.

However, we can strengthen the result if we consider other propositions; if the agent's credences are sufficiently rich, we can show that $Cr_{Current} = Cr_{Best}$.

Let's assume that, for you, $Cr_{Current}(Rain) = 0.8$, and $Cr_{Current}(\sim Rain) = 0.2$. Now find some proposition Q, such that $Cr_{Current}(Q) = 0.7$. And suppose you are considering the following question:

Question 2: Rain, or Q?

Since $Cr_{Current}$(Rain) is greater than $Cr_{Current}$(Q), you regard Rain as a better guess. This means should choose Cr_{Best} so that it will guess Rain as well. But this constrains your credences more tightly than the previous question did: instead of just choosing so that Cr_{Best}(Rain) > Cr_{Best}(~Rain), you need to make sure that Cr_{Best}(Rain) > 0.7.

Considering questions like Question 2 is what we need to get Immodesty, if we make some additional assumptions. Namely, we must assume that your credences *and* the credences you regard as best are defined over a very rich space of possibilities, so that for every real number *n* we can find a proposition A for which $Cr_{Current}(A) = Cr_{Best}(A) = n$. (This is a very strong assumption! I'll return to it in a minute.)

With that assumption in hand, suppose for reductio that you choose a Cr_{Best} that differs from $Cr_{Current}$ in some respect. Hold all of the values of $Cr_{Current}$ fixed except for its assignment to one proposition P, and without loss of generality, assume that $Cr_{Best}(P) < Cr_{Current}(P)$. In that case, there is at least one question that $Cr_{Current}$ expects Cr_{Best} to answer falsely. To see why, find some proposition R such that $Cr_{Best}(P) < Cr_{Current}(R) = Cr_{Best}(R) < Cr_{Current}(P)$. In response to the question: {P, R}, $Cr_{Current}$ will regard P as the better answer, accuracy-wise. But Cr_{Best} will license guessing R. So Cr_{Best} will license the wrong answer, by the lights of $Cr_{Current}$. This means that Cr_{Best} is not the credence function that $Cr_{Current}$ recommends, after all; $Cr_{Current}$ will not recommend it if it licenses bad guesses. This contradicts our assumption.

The only credence function that will license all and only *good* guesses, by the lights of $Cr_{Current}$, is $Cr_{Current}$ itself. So $Cr_{Current}$ is uniquely self-recommending.

The "personal" argument for Immodesty is now on the table. However, it does rely on a controversial richness assumption. In the next subsection I will look at the possibilities for depersonalizing the argument, which will underscore how crucial this assumption is.

2.3 Immodesty: the impersonal approach

Is there a way to depersonalize that argument for Immodesty, as we did with Truth-Directedness?

In the case of Truth-Directedness, we showed that a higher credence in a true proposition is more accurate than a lower credence by looking at how each would fare in different credal environments; we saw that the higher credence weakly dominates the lower credence in terms of licensing true guesses.

What would a parallel argument for Immodesty look like? We would have to show that a 0.8 credence in P, say, sees itself as better than a credence of 0.6 in P, but also as better than a credence of 1 — regardless of the rest of the credence function in which it is embedded.

How could we do this in terms of guessing? Unfortunately, the most straightforward route that one might try doesn't work. Consider a 0.8 credence in P, and some arbitrary proposition Q which has a different truth value from P. Now assign some arbitrary credence to Q, assuming a uniform distribution over [0,1].

It's clear that 0.8 credence in P doesn't *dominate* either 0.6 or 1, because there are possible situations in which one of those credences does better than 0.8. (For example, suppose P is true, Q is false, and cr(Q) = 0.9. In that case, it would be better to have credence 1 in P, since this would license you to guess truly, whereas credence 0.8 would license you to guess falsely.)

It's also not the case that 0.8 does better in expectation than other values. Consider the following four possibilities, which are all of those in which you will give a unique answer to {P, Q} that counts for or against your accuracy:

w1	P is true, Q is false, and cr(Q) < cr(P)
w2	P is true, Q is false, and cr(Q) > cr(P)
w3	P is false, Q is true, and cr(Q) < cr(P)
w4	P is false, Q is true, and cr(Q) > cr(P)

In w1 and w4, you guess truly. In w2 and w3, you guess falsely. The probability of w1, from the point of view of 0.8 credence in P, is 0.8 * 0.8 = 0.64, and the probability of w4 is 0.2 * 0.2 = 0.04. So, the probability of guessing truly is 0.68, and the probability of guessing falsely is 0.32.

Unfortunately, from the point of view of a 0.8 credence in P, you could *improve* your prospects for guessing truly if you became even more confident in P. Consider a 0.9 credence in P. From the point of view of a 0.8 credence in P, the probability of a 0.9 credence guessing truly is 0.73, and the probability of guessing falsely is 0.27. (Pr(w1) = 0.8 * 0.9 = 0.72, and Pr(w4) = 0.2 * 0.1 = 0.01; add these to get the probability of guessing truly.) So it look like a good idea to move from 0.8 to 0.9.

This means it does not work to argue that a credence of 0.8 is immodest in isolation: a 0.8 credence does not expect itself to be best if it is dropped into an arbitrary credence function. This argument seems to show that it's not possible, using the guessing view, to give a depersonalized version of Immodesty.

The reason the depersonalized approach does not work for Immodesty goes hand in hand with the fact that it *does* work for Truth-Directedness. In isolation, a 0.8 credence in P does not license any guesses beyond {P, ~P}. (Even that contentiously assumes that 0.8 credence in P comes with a 0.2 credence in ~P. One could argue that 0.8 credence in P does not license any guesses at all.) So from the point of view of that single credence — if that is all we are holding fixed — the only thing we can try to ensure is a true guess on that question. Of course, this means that it would be better to have higher credence in P; as a corollary, 0.8 credence in P is not immodest on its own.

Notice that as soon as we add any credence in Q, the calculation changes: we will be able to defend Immodesty for the comparative relationship between P and Q. Because agents are immodest about their *comparative* judgments, it will turn out that as an agent becomes more opinionated, they will also become more immodest. Additionally, if there is another way to rationally fix some of an agent's numerical credences (such as the Principal Principle), this will constrain their other credences numerically. Suppose Tim and Jim each become opinionated about just one more proposition: out of a ten-ticket lottery, they both become 0.7 confident in the proposition Ticket, that the winning ticket will be one of tickets #1-7. Now Tim will consider his 0.8 credence in Rain to be better than Jim's 0.6 credence in Rain, and vice versa: although they expect one another to guess truly on {Rain, ~Rain}, each will prefer to give *his own* answer to {Rain, Ticket}. Furthermore, Tim will prefer to have his credence in Rain be greater than 0.7. Jim will prefer his credence in Rain to be between 0.5 and 0.7.

Still, the bottom line is that we can't get full Immodesty without full richness. I will look at that assumption more directly in the next subsection.

2.4. Richness

I have argued that we can only defend Immodesty if we assume richness. Agents who have gaps in their credences, like Tim and Jim, won't turn out to be immodest. And we can't defend a "depersonalized" version of Immodesty like we could for Truth-Directedness: credences are only immodest in the context of a rich and fully specified credence function.

Although the richness assumption is controversial, similar assumptions are widespread in this literature. Gibbard's [2008] argument for a practical understanding of Immodesty requires us to respond differently to bets whose payouts differ infinitesimally, which would require us to discriminate between these bets in our credences as well. Traditional accounts that define credences in terms of betting behavior similarly assume that an agent's preferences must be extremely rich and specific. For instance, Ramsey's representation theorem assumes that an agent must be able to make specific judgments about "gambles" in order to allow the theorist to differentiate her credences in different propositions.[16] Both of these approaches employ a similar assumption to the one I make here: that any time we find a gap in an agent's opinions, we can always "fill it in" with opinions about other propositions.[17]

"Comparativist" views of credence, which ground credences in comparative confidence judgments, have the result that our credences are only numerically precise if our confidence ordering is "superfine".[18] If there are any gaps, comparativists would represent our doxastic state with "mushy", or imprecise credences. According to comparativism, then, we are wrong to say that Tim and Jim have different precise credences in Rain. (Presented with two simple agents who look a bit like Tim and Jim, a comparativist would say that they each have a single comparative judgment: $\sim$Rain $\prec$ Rain. If we wanted to describe this state numerically, we would use the same interval, $(0.5, 1]$, to characterize both Tim and Jim's opinions about Rain.) A comparativist would not be surprised to find that Tim does not regard his credences are more accurate than Jim's — after all, Tim and Jim's doxastic states are the same.

Indeed, formal accounts of accuracy do not usually deliver immodesty for imprecise credences.[19] Both standard epistemic utility theory and the guessing account of accuracy have the result that an agent with imprecise credences regards those imprecise credences as no better or worse than certain precise states. Standard epistemic utility theory also has the result that an agent with *precise* credences regards her own doxastic state as no better or worse than certain *imprecise* states. (I will return to this topic in Chapter 8.) So if we think that agents like Tim and Jim are better represented as imprecise than as "sparse" or "gappy", we should perhaps not expect them to be immodest.

Adding up these points, we can make two observations about the status of the richness assumption. First, given its prevalence in the literature, it is unsurprising that we should need it here as well. And second, on at least one view of what credences *are*, we should predict that a richness assumption is needed to defend Immodesty in particular. (Note however that this still leaves us in a weaker position than epistemic utility theory: even if EUT cannot defend Immodesty for imprecise credences, it can still defend it for Tim and Jim, assuming they do have precise credences.)[20]

Richness assumptions can more easily slip in unnoticed (or alternatively: can more easily be justified!) in the context of betting-based arguments. In that setting, the monetary payouts of different bets give us a natural way to fill in a very rich set of propositions, over which it is

---

[16] Ramsey [1931]. See Titelbaum [2022], ch. 8, for an overview of different representation theorems.
[17] See also Builes, Horowitz and Schoenfield [2022] for further discussion of the richness assumption.
[18] See Konek [2019]; see Stefánsson [2017] for a defense of comparativism. See Builes, Horowitz and Schoenfield [2022], and Chapter 8 of this book, for further discussion.
[19] See Seidenfeld et al. [2012], Mayo-Wilson and Wheeler [2016], Schoenfield [2017], Berger and Das [2020], and Konek (forthcoming). Konek is the exception to the general claim here — that imprecise credences are modest.
[20] Thanks to Alejandro Pérez Carballo for helpful discussion on the dialectic in this section.

natural to assume that agents have preferences. But with the guessing approach, the "payouts" (if we can consider the upshots of guessing truly and falsely as payouts) are too coarse-grained. We need to find richness elsewhere. What we need is a set of propositions whose *comparisons* can yield those same kinds of fine-grained distinctions; this is difficult, if we want to account for gappy agents like Tim and Jim.

We do have an additional option in this case, however, which was not open to us in the case of Truth-Directedness. Immodesty is supposed to be a constraint on *rational* credences. So if we can argue that rational credences must obey this richness constraint, we have a good response to the objection. (This argument would not work for Truth-Directedness, because Truth-Directedness is supposed to apply regardless of rationality.)

It is not entirely implausible that this could be a rational requirement. Suppose that it's a rational requirement to match your credences to known chances, as the Principal Principle recommends. And suppose you are rationally required to regard the following situation as possible: a coin factory is able to make coins at any possible weighting. In countenancing this possibility, you would obey the richness requirement. All of these possible coins figure into propositions to which you assign different values: in fact, for any two propositions P and Q to which you assign different credences, we can find some possible coin such that your credence that the coin lands heads will lie between your credences in P and Q. The proposition that this coin lands heads can be slotted in for proposition A in the argument in section 2.2.[21]

2.5 What Immodesty means

I would like to close this section by returning briefly to Immodesty, and clarifying how it should be interpreted. I began with Lewis's example of the product magazines, and other examples where we imagine a person having a certain kind of self-reflective or introspective thought: are my beliefs the most accurate? Could I do better by adopting other ones? My argument, which involved imagining oneself preparing to take a quiz, had the same introspective quality.

These are helpful illustrations because it's clear that if someone were to explicitly hold a modest attitude ("I'd be more accurate if I believed that it's raining right now, rather than believing that it's not raining") they would look clearly irrational and incoherent. But Immodesty is not a requirement on metacognition. There is no rational requirement to *have the belief* (or credence) that one's beliefs are accurate — at least if this involves anything like the kind of conscious thought that we appealed to when defending and illustrating Immodesty. Rather, Immodesty is a kind of coherence between a belief state and a method of assessing accuracy. If you end up with a picture according to which rational credences are *not* immodest, you have made a mistake: either you are not really describing rational credences (or credences at all!), or you are not really describing accuracy. Rational credences and accuracy must stand in this sort of relation to one another as part of their functional roles.

But if Immodesty is not a metacognitive requirement, it is fair to ask: what is the role of the argument I have given in favor of it, which asks us to consider a believer reflecting on her own credences? I propose that we see that argument as a test. If accuracy and rational credence

---

[21] If, for some reason, you are skeptical of deference to numerical chances, we could plausibly set up Question 2 without making that assumption. What we need to get Question 2 off the ground is some proposition P such that $Cr_{Best}(Rain) < Cr_{Current}(P) < Cr_{Current}(Rain)$. Savage ([1954], pp. 38-39) argues that any time we are more confident of P than of Q, we can find some intermediate proposition R (such that $Q \prec R \prec P$) by taking the disjunction of Q and some unlikely series of coin flips, for some possible coin. As long as we can find that proposition R, we can set up Question 2 by asking you to guess between P and R.

are related in the right sort of way, they will survive being put together by a self-conscious reflective agent. The fact that the guessing account survives the test shows that it can support Immodesty.

Seeing this argument as a test also points to another possibility for understanding the role of richness. We might see the richness assumption as providing the appropriate environment to test for Immodesty: one might think that things have gone especially wrong if an agent with incredibly specific and fine-grained opinions would be willing to scramble them with no new evidence.

3. Conclusion

Immodesty and Truth-Directedness are both necessary, I have argued, to show that we are talking about *accuracy* rather than something else. I have shown here that the guessing account can make sense of both principles, albeit in limited forms. For Truth-Directedness, we were able to support a comparative principle, as well as a weak dominance principle:

> **Comparative Truth-Directedness:** For any true proposition P and false proposition Q, it is always more accurate to have Cr(P) > Cr(Q) than to have Cr(P) ≤ Cr(Q).
>
> **Weak Truth-Directed Dominance:** For any true proposition P, and two credences a and b such that a > b, a is more accurate than b in the following sense:
>
> For any possible credence n in a false proposition, Q, the combination of a and n licenses at least as many true guesses as the combination of b and n, and for some values of n, the combination of a and n licenses strictly more true guesses than the combination of b and n.

For Immodesty, we were able to support the following:

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

By making sense of these two essential thoughts, this chapter puts the guessing framework on the table as a theory of accuracy. It also highlights some essential limitations of the guessing framework. It turns out that we cannot fully support Non-Comparative Truth-Directedness in the same form that EUT can. EUT, however, builds this principle in from the ground floor, either as an axiom or as a consequence of strict propriety. If you are wedded to Non-Comparative Truth-Directedness as being definitional of accuracy, this may be a reason to prefer EUT to the guessing framework. But to those who are more open-minded about the issue, the guessing framework invites us to replace Non-Comparative Truth-Directedness with some similar, though weaker, relatives.

It also turned out that we can only support Immodesty with the help of a controversial richness assumption. This assumption will come up again in this book: it is often needed to bridge the gap between the comparative notion at the heart of guessing and the precise numerical values which we customarily use to represent credences.

The next several chapters investigate how accuracy, as the guessing account conceives it, figures into arguments for various rational requirements.

(Abridged) Bibliography

Berger, Dominik, and Das, Nilanjan. [2020] "Accuracy and Credal Imprecision." *Noûs* 54(3): 666-703.

Berker, Selim. [2013a] "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122 (3):337-393.

— [2013b] "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23(1): 363-387.

Builes, Horowitz, and Schoenfield [2020] "Dilating and Contracting Arbitrarily". *Noûs* 56(1): 3-20. https://doi.org/10.1111/nous.12338

Campbell-Moore, Catrin, and Levinstein, Benjamin [2021]. "Strict Propriety is Weak". *Analysis* 81(1): 8-13.

Carter, J. Adam. [2020] "Sosa on Knowledge, Judgment, and Guessing." *Synthese* 197: 5117-5136.

Christensen, David. [1996] "Dutch-book Arguments Depragmatized: Epistemic Consistency for Partial Believers." *Journal of Philosophy* 93(9): 450-479.

Cohen, L. Jonathan. [1974] "Guessing." *Proceedings of the Aristotelian Society*, New Series, Vol. 74 (1973-1974), pp. 189-210.

Daoust, Marc-Kevin. [2020] "Should Agents be Immodest?" *Analytic Philosophy* 62 (3):235-251

Dartnall, Terry. [1984] "Knowing and Non-Accidental Guessing." *Analysis* 44(1): 38-41.

Dorst, Kevin and Mandelkern, Matthew. [2021] "Good Guesses". Philosophy and Phenomenological Research 105 (3):581-618.

— [2022] "Assertion is weak". *Philosophers' Imprint* 22.

Easwaran, Kenny. [2016] "Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities". *Noûs* 50 (4):816-853.

Fitelson, Branden and McCarthy, David. [2014] "Toward an Epistemic Foundation for Comparative Confidence," ms. Available at: http://fitelson.org/cc_handout.pdf.

Fitelson, Branden, with Easwaran, Kenny, and McCarthy, David. [2016] *Coherence*, ms. Available at: http://fitelson.org/coherence/coherence_duke.pdf

Gendin, Sidney. [1998] "Guessing." *Philosophy* 26, 435-440.

Gibbard, Allan. [2008] "Rational Credence and the Value of Truth." Gendler and Hawthorne, eds., Oxford Studies in Epistemology 2. Oxford University Press.

Goodman, Jeremy and Holguín, Ben. [2022] "Thinking and Being Sure." *Philosophy and Phenomenological Research* 106 (3):634-654. doi: 10.1111/phpr.12876

Hawthorne, John; Rothschild, Daniel & Spectre, Levi [2016]. Belief is weak. Philosophical Studies 173 (5):1393-1404.

Holguín, Ben. [2022] "Thinking, Guessing, and Believing." *Philosophers' Imprint* 22 (1):1-34.

Holton, Richard. [2013] "Intention as a Model for Belief" in Manuel Vargas and Gideon Yaffe (eds.) Rational and Social Agency: Essays on the Philosophy of Michael Bratman (Oxford: Oxford University Press) pp. 12-37.

Horowitz, Sophie. [2021] "Shrinking three arguments for conditionalization." *Philosophical Perspectives* 35(1): 303-319.

— [2019] "Accuracy and Educated Guesses." Gendler and Hawthorne, eds., *Oxford Studies in Epistemology* 6. Oxford University Press.

Jeffrey, Richard [1986] "Probabilism and Induction." *Topoi* 5: 51-58.

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. pp. 263-297.

Kerbel, Gabrielle. [ms a] "Truth-Directed Guessing."

— [ms b] "Guessing with Imprecise Credences."

— [ms c] "Guessing on Principle."

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

Lewis, David. [1971] "Immodest Inductive Methods" *Philosophy of Science* 38 (1):54-63.

Lindley, Dennis. [1982]. "Scoring Rules and the Inevitability of Probability." *International Statistical Review / Revue Internationale de Statistique* 50(1): 1-11.

Maher, Patrick. [2002] “Joyce’s Argument for Probabilism.” *Philosophy of Science* 69(1): pp. 73-81.

— [1993] *Betting on Theories*. Cambridge University Press.

Montague, Phillip [1981]. “Self-defense and choosing between lives”. *Philosophical Studies* 40 (2):207 - 219.

Radford, Colin. “Knowledge — by Examples”. *Analysis* (27)1: 1-11.

Ramsey, Frank [1931] Truth and Probability. In: *The Foundations of Mathematics and other Logic Essays*. Ed. by R. B. Braithwaite. New York: Harcourt, Brace and Company, pp. 156–198.

Schervish, Mark; Seidenfeld, Teddy; and Kadane, Joseph. [2009] “Proper Scoring Rules, Dominated Forecasts, and Coherence”. *Decision Analysis* 6(4):202-221. https://doi.org/10.1287/deca.1090.0153

Schoenfield, Miriam. [2017b] “The accuracy and rationality of imprecise credences”. *Noûs*, *51*(4), 667–685.

Schroeder, Mark. [2017] “The Epistemic Consequences of Forced Choice”. *Logos and Episteme* 8 (3):365-374.

Seidenfeld, Teddy; Schervish, Mark; and Kadane, Joseph. [2012] “Forecasting with Imprecise Probabilities.” *International Journal of Approximate Reasoning* 53:1248–61.

Sorensen, Roy. [1984] “*Unique alternative guessing*.” Logique Et Analyse 27 (5):77.

— [1982] “Knowing, Believing, and Guessing”. *Analysis* 42(4): 212-213.

Sosa, Ernest. [2015] *Judgment and Agency*. Oxford: Oxford University Press.

Stefánsson, H. O. [2016] “What is “Real” in Probabilism?” *Australasian Journal of Philosophy* 95 (3):573-587.

Titelbaum, Michael. [2022] *Fundamentals of Bayesian Epistemology*. Oxford University Press.

Weatherson, Brian. [2011] “Accuracy Measures on Conditional Probabilities.” Blog post. Available at: http://tar.weatherson.org/2011/04/25/accuracy-measures-on-conditional-probabilities/