

Dilating and Contracting Arbitrarily

David Builes, Sophie Horowitz, and Miriam Schoenfield

Forthcoming in *Noûs*

Abstract

Standard accuracy-based approaches to imprecise credences have the consequence that it is rational to move between precise and imprecise credences arbitrarily, without gaining any new evidence. Building on the Educated Guessing Framework of Horowitz (2019), we develop an alternative accuracy-based approach to imprecise credences that does not have this shortcoming. We argue that it is always irrational to move from a precise state to an imprecise state arbitrarily, however it can be rational to move from an imprecise state to a precise state arbitrarily.

1. Introduction

Suppose you have no idea whether *P*. You're completely clueless. Can you rationally move from a state of uncertainty about *P* to a state in which you're opinionated about *P*, without a change in your evidence? What about the reverse? Suppose you start out with a pretty firm opinion about *P*. Can you rationally move from your opinionated state to a state of uncertainty just for kicks? Many people have the intuition that it's always (epistemically) irrational to revise one's beliefs without a change in one's evidence. But why would this be? Plausibly, it's fine for our preferences to shift for arbitrary reasons. Why not our belief states?

There are many ways to address the question of what would make arbitrary shifts in one's opinions irrational, and which answers will be satisfying will depend on what general epistemological framework one is working within. Here we'll be thinking about the question from an accuracy-based perspective. In other words, we'll be assuming that the requirements of epistemic rationality are grounded in a concern with being accurate which, in the broadest sense, can be thought of as a concern with "getting things right" – having one's belief state in some sense match, or approximate, the way the world really is.

If we assume that rational agents' belief states are representable by precise probability functions, then there is a straightforward answer to the question of why a rational agent won't shift belief states arbitrarily. According to the most popular ways of thinking about credal accuracy, an agent with a certain probability function will regard any alternative probability function as worse, from

the point of view of accuracy, than her own.¹ Shifting from one probability function to another without new evidence will therefore look like a bad idea from the perspective of an agent who wants her belief state to be accurate.

However, it's plausible that in some cases – perhaps cases in which our evidence is ‘incomplete’ or ‘non-specific’ – we lack any precise credence (consider, for example, the proposition that one of the authors of this paper is currently wearing a striped shirt). Consideration of such cases has led some philosophers to argue that we sometimes adopt and/or ought to adopt attitudes that are ‘imprecise’ and hence better represented by a set of credence functions (a ‘representor’) rather than a single one (more on how these sets are generated later).² But once imprecise credal states enter the mix, the story about whether accuracy considerations permit arbitrary doxastic shifts becomes significantly more complicated. This story is the topic of our paper.

To get clearer about the question we'll be addressing, it will be helpful to introduce some terminology. We'll think of precise states as represented by either single credence functions or, when convenient, sets containing a single credence function. Imprecise states are represented by sets containing more than one credence function.³ We'll say that an agent *dilates* when she moves (without a change in evidence) from a precise state, p , to an imprecise state M that has p as a member. We'll say that an agent *contracts* when she moves (without a change in evidence) from an imprecise M to a precise p which is a member of M .

Our question is the following:

Question: Are there accuracy-based reasons to avoid dilating? Are there accuracy-based reasons to avoid contracting?

We'll be arguing for:

Answer: There are accuracy-based reasons to avoid dilating, but there are no accuracy-based reasons to avoid contracting.

In defending our answer we'll be assuming a particular view about what having an imprecise credal state amounts to. This view is based on *comparativism* about subjective probability more broadly: the view that we should understand subjective probabilities as representations of facts concerning the agent's comparative confidence ordering. We'll be thinking of agents with imprecise probabilities as those whose comparative confidence ordering is incomplete (more details on this picture later).

Here's the plan: We'll first explain why the standard epistemic utility theory framework (EUT) permits both dilation and contraction, and why this result is problematic. We'll then show that an alternative to EUT, Sophie Horowitz's educated-guess framework, forbids dilation. This is a point in its favor. However, as we'll see, like EUT, the educated guess framework permits

¹ For a discussion of this feature in the context of epistemic utility theory see Pettigrew (2016). This feature is also present in Horowitz's (2019) educated guess framework.

² Some defenses of imprecise credences include Levi (1985), Jeffrey (1983), Joyce (2005, 2009, 2010), Sturgeon (2008), and Kaplan (2009).

³ We will sometimes abuse terminology and talk as if the state *is* a set.

contraction. We will explain why the permissibility of contraction makes sense, at least from an accuracy perspective, and we'll conclude by discussing the epistemological significance of the asymmetry between dilating and contracting that we're defending in this paper.

2. EUT Permits Dilating

In recent literature, accuracy considerations have been brought to bear on questions about imprecise credences using *epistemic utility theory* (EUT).⁴ The central tool in EUT is that of a *scoring rule*, which intuitively measures the “distance” between one’s credences and the truth at any given possible world. Scoring rules are also interpreted as measures of *epistemic value*: higher credences in truths are better than lower credences in truths, and lower credences in falsehoods are better than higher credences in falsehoods. By using an appropriate scoring rule, one can give all sorts of interesting decision-theoretic arguments for different rational requirements.

EUT was developed in the context of precise probabilities. How should it be extended to account for imprecise probabilities?

There are two ways one might measure the accuracy of an imprecise credal state $\{c_1 \dots c_n\}$ on EUT. First, one might assign an imprecise credal state (at a given world) a number, representing its accuracy at that world: let's call that the “numerical approach”. Alternatively, one might assign an imprecise credal state some non-numerical score, for example a set of numbers: let's call that the “non-numerical approach”. A natural way to spell out the non-numerical approach would be to use the set $\{A(c_1) \dots A(c_n)\}$, where $c_1 \dots c_n$ are the precise credence functions in the agent's representor, and $A(c_1) \dots A(c_n)$ are the accuracy scores of those precise credence functions at a given world.

In this section we'll explain why we think both the numerical and non-numerical approaches permit dilation, and why we take this to be a problematic result.

Let's start with the numerical approach. Results discussed in Mayo-Wilson and Wheeler (2016), building on Seidenfeld et al (2012), show that given certain plausible constraints on an accuracy measure, for any precise state, p , there is an imprecise state M , such that, in every world, M and p are equally accurate.⁵ This suggests that from an accuracy perspective, if you are in state p , there is no reason not to dilate to M . After all, if you dilate to M you're guaranteed to do just as well

⁴ For a comprehensive overview of EUT, see Pettigrew (2016). For the implications of of EUT on questions concerning imprecise credences see Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2016), Schoenfield (2017), Berger and Das (2020), and Konek (forthcoming).

⁵ The three relevant constraints are continuity, extensionality (the accuracy score is a function only of the credal state and the proposition's truth value), and admissibility (probabilistic states are not dominated). Note that Konek (forthcoming) denies admissibility. On the accuracy measures he discusses probabilistic belief states are sometimes accuracy dominated. Agents who value accuracy in the way that motivates imprecision thus ought to sometimes violate the Principal Principle. Since, as will become clear, one of our desiderata is to make the accuracy framework compatible with the Principal Principle, we will be setting Konek's view aside for the remainder of the paper. However, it is worth noting that Konek's view is one which forbids both dilation and contraction, at least as long as the agent doesn't change her epistemic values.

as you would by staying at p , no matter what the world is like. So the numerical approach can't motivate a prohibition on dilating. It also can't motivate a prohibition on contraction, for a similar reason: for any imprecise state, there's a precise state that's guaranteed to be equally accurate.

Can the non-numerical approach do better? We think not. If accuracy scores are given by non-numerical objects, like sets, for example, then to figure out whether dilating is permitted, we need a more complicated account of how to compare the accuracy of different belief states.

Here's one proposal: extending the supervaluationist picture that infuses much of the literature on imprecise credences, we could say that the accuracy of A is less than the accuracy of B if and only if, for every a_i in A and b_j in B , the accuracy of a_i is less than the accuracy of b_j . On such a picture, there will never be accuracy-based reason to prefer adopting a precise credence function p over an imprecise state M , which contains p , or vice versa. For if M contains p , then there is some credence function in M (namely p !) which is no less accurate than p in every world. It follows, then, that if you have p , there is no accuracy-based reason to avoid dilating to some M that contains p as a member.⁶ (There is also no reason, if you start off in M , to avoid contracting to p .) There are certainly other ways to think about and compare accuracy-scores if accuracy is measured non-numerically. But we are doubtful that other ways of measuring can solve the present problem; as Schoenfield (2017, "Imprecision-2") and Berger and Das (2020, Proposition 4) show, for any such measure satisfying certain plausible constraints, at least some instances of dilating will be permitted.⁷

The supervaluationist, non-numerical method leads to especially worrisome consequences because of the following feature: on this approach, the accuracy of an imprecise credence is *no better or worse* than the accuracy of *any* of the precise probability functions it contains. This seems to permit believers to jump around between *precise* credal states willy-nilly. Consider:

Delia: Delia's credence that there is intelligent life on other planets is only 0.1. (This is supported by her evidence as well). But she thinks: "it sure would be *fun* to believe that there was life on other planets!" Alas, as somebody highly committed to rationality and accuracy she knows she can't just hop from one precise credence to another... But then Delia remembers that she has another option: imprecise credences! Consulting some results from epistemic utility theory, Delia decides that there is no accuracy-based reason not to adopt $[0,1]$. So she does. Then, consulting epistemic utility theory again, she decides that there is no accuracy-based reason not to adopt a precise credence. She adopts a precise credence of 0.9 that there is intelligent alien life and enjoys her new belief state very much.

⁶ Does it help to suggest that M and p are incomparable with respect to accuracy? We think not. We agree with Schoenfield (2017, note 18) and Berger and Das (2020, p.21 section 6) who argue that if one claims that two states are incomparable with respect to accuracy in every world, then there can be nothing, accuracy-wise favoring one over the other.

⁷ The constraints here are just admissibility and extensionality. What Berger and Das show is that if M assigns the same set of numbers to every cell of an n -celled partition, and p assigns $1/n$ to each cell of an n -celled partition, then p will be no less accurate than M in every world. It follows from this result that on any non-numerical measure satisfying admissibility and extensionality, even if we know, say, that a coin is fair, there will be no reason for preferring 0.5 to, for example, $[0,1]$ in the proposition that the coin lands Heads. Seidenfeld et al (2012) propose a non-numerical score that forbids dilating but it does not satisfy the extensionality constraint. See also 29.

On the non-numerical approach we've been sketching, each of these transitions is rational. But it's preposterous to think that one can rationally move from 0.1 to 0.9 without a change in evidence, simply by stopping at $[0,1]$ along the way!⁸

There is also a more general, unattractive consequence of *any* theory that permits dilating. That is, if dilating is permitted, there can be no requirement to conform to the Principal Principle. If we only care about accuracy, and if $[0, 1]$ is no less accurate than 0.5 in every world, how could we be required to adopt 0.5 credence that a fair coin will land heads, rather than $[0,1]$?⁹

The Principal Principle is one of the least controversial principles governing the rationality of credences, and Delia's reasoning seems like a paradigmatic instance of irrationality. But no version of EUT that meets widely accepted plausible constraints on an accuracy measure can explain why an agent interested in accuracy would want to maintain the credences recommended by the Principal Principle, and some otherwise attractive elaborations of EUT will vindicate Delia's reasoning. The source of both of these problems is that EUT permits dilation. We take these considerations to pose a serious challenge to EUT, and to motivate considering an alternative framework for thinking about the accuracy of imprecise credal states.

In what follows, we will introduce a new way of assessing the accuracy of imprecise credences, building on the guessing-based framework for accuracy developed in Horowitz (2019). According to this view, credences are accurate or inaccurate in virtue of the all-or-nothing 'guesses' that they license. Using this simple and intuitive framework, Horowitz argues that one's credences should satisfy the axioms of probability, given some very natural norms governing educated guesses. Horowitz also derives a version of *Immodesty*, which is roughly the claim that a rational agent should expect her own (precise) credences do better accuracy-wise than any other (precise) credences.¹⁰

A major advantage of this approach, over EUT, is that it does not require us to attach accuracy values to credences, or to aggregate the values of an agent's credences in different propositions. In other words, EGT – unlike EUT – does not include a scoring rule. Although, like EUT, Horowitz's framework is based on the thought that norms of rationality can be explained and motivated by the aim of accuracy, that aim is not captured by thinking of belief states as a whole

⁸ The supervaluationist, non-numerical method does not *automatically* commit one to the consequence that transitions like Delia's are rational: one could supplement the view with a diachronic norm that would rule such transitions out. We are skeptical that such a norm could be well-motivated but we will not pursue the issue here. Thanks to Chris Meacham for pressing us on this point.

⁹ See Levinstein (2019) and Schoenfield (2017) for further discussion of this problem. While this point is usually described as EUT conflicting with the Principal Principle, some have argued that the Principal Principle should simply be understood as an implicit *definition* of our concept of "chance". On this way of understanding the Principal Principle, the conflict is more perspicuously described as EUT entailing that, contrary to our best science, there are no precise chances. See Builes and Spencer (MS) for a recent defense of this point of view.

¹⁰ These derivations of Probabilism and Immodesty have important advantages over the standard derivations given by EUT. For example, using EUT, one can only derive these consequences given substantive constraints on allowable scoring rules. However, some philosophers (e.g. Maher (2002), Gibbard (2008)) disagree with these substantive constraints. We'll be assuming for the remainder of the paper that the credence functions in agents' representors are probabilistic. Horowitz doesn't discuss imprecise credences in her paper, but our generalization of her framework will apply to both precise and imprecise credences, and will preserve the result that precise credences are "strictly self-recommending": they regard themselves as strictly more accurate than any imprecise credal state as well as over any other precise credal state.

as scoring better or worse on some accuracy scale. EUT, on the other hand, standardly *does* work with aggregate scores of an agent’s entire belief system. This aspect of the view is what seems to justify unsavory “tradeoffs” between true and false beliefs in different propositions. Some epistemologists have rejected EUT for this reason, and have taken this objection to EUT as a reason to reject the aim of accuracy altogether. However, if the real problem is aggregation, these epistemologists should feel free to embrace accuracy-first epistemology with EGT instead. Furthermore, the scoring rules commonly used in EUT build in controversial assumptions that some epistemologists deny.¹¹ So avoiding scoring rules carries some significant advantages for EGT. In our expansion of EGT, we will preserve these advantages, as well as developing one more: its treatment of dilation.

We’ll first show how the guessing framework can be generalized to apply to imprecise credences. We will then argue that, unlike EUT, the educated guessing theory (hereafter, for symmetry’s sake, “EGT”) can explain the irrationality of dilation. Finally, we will go on to argue that, perhaps surprisingly, there are important differences between dilation and contraction. While there are good accuracy-based reasons to avoid dilation, there is no accuracy-based reason to avoid contraction.

3. Educated Guessing and Accuracy

First, some background on EGT: following Horowitz, we assess the accuracy of credences by looking at the *educated guesses* that our credences license. We can think of educated guesses as answers to forced choice questions, perhaps given various suppositions. An agent does well with respect to a guess if her guess is true, and poorly if her guess is false. (“Well” and “poorly” are not defined in numerical terms, here; using this approach, we are just meant to think of ourselves as desiring, for each question we might encounter, that we answer it correctly.) As one might imagine, different credal states license different guesses: If an agent is more confident in P than Q, for example, and asked to guess either P or Q, then she is licensed to guess P (and not licensed to guess Q).¹²

Forced choice questions can also be posed under various suppositions. One might be asked: Supposing it rains tomorrow, guess between: (A) Ali will bring an umbrella, and (B) Ben will wear rain boots. In response to this kind of question, if an agent is more confident of A than B conditional on its raining tomorrow, she is licensed to guess A (and not licensed to guess B). If she is not more confident of either A or B conditional on its raining tomorrow, she is in a state that licenses either guess.

Since our guesses can be true or false, our credences can get things right or wrong by licensing true or false guesses. So, we can understand an agent’s credences as being accurate insofar as the

¹¹ In particular, Maher (2002) and Gibbard (2008) are skeptical of the constraints on scoring rules that standard EUT assumes. Berker (2013) develops the problem of epistemic tradeoffs, which educated guessing avoids. See Horowitz (2019), section 4.1, for further discussion of the comparison between educated guessing and EUT.

¹² Horowitz (2019) does not give a general definition of the notion of licensing. We will understand it as follows: A doxastic state S licenses a guess G so long as it is compatible with maximal rationality that, for the purpose of doing well on the relevant guessing question, an agent in state S guesses G.

guesses they license are true, and inaccurate insofar as the guesses they license are false. (Following Horowitz, we will assess the accuracy of suppositional guesses only if the supposition is true; otherwise, their accuracy is undefined.)

4. Educated Guesses for Imprecise Credences

To use the guessing framework to motivate norms about imprecise credences, we need to first give an account of which guesses are licensed when our credences are imprecise. That's the aim of this section. (Precise credences are easy: if you are asked to guess between various options, you are licensed to guess exactly those options in which your credence is highest.) Which guesses are licensed if you are in an imprecise doxastic state? To answer that, it will be helpful to first step back and ask what an imprecise doxastic state is. All we have said so far is that an imprecise doxastic state is one in which it makes sense to ascribe a set of probability functions to you – a representor – rather than just a single probability function. But what is it about you that determines the need for more than one probability function in your representor – and what determines *which* probability functions are included? Our view is based on *comparativism* about subjective probability.¹³

Comparativism is based on the intuitive thought that while numerical probabilities *represent* belief states, there's nothing about our belief states that mandates a unique numerical representation. In other words, there's nothing "0.69-ish" about my degree of confidence in P, beyond the fact that .69 can serve as an adequate representation of my degree of confidence within a particular representational system. But 69, for example, or 732.6 for that matter, would work just as well, provided the system was structured in the right way. In other words, the comparativist thinks that what makes it the case that I have a credence of 0.69 in P is really a *structural* fact about how my degree of confidence in P is related to my degrees of confidence in other propositions.¹⁴

In this spirit, we hold that the probability functions in an agent's representor are all and only the ones that are compatible (in the sense defined below) with the agent's comparative confidence judgments.¹⁵ So if, for example, an agent is more confident in P than she is in Q, every probability function in her representor will assign greater credence to P than to Q. If she is more confident in P given Q than she is in P given $\sim Q$, every probability function in her representor will assign greater conditional credence to $(P|Q)$ than to $(P|\sim Q)$. We will state the view in its most general form in terms of conditional comparative confidence judgments; we understand unconditional comparative confidence judgments as just one's conditional comparative confidence judgments given the tautology.

¹³ For alternative explications of what is required for a credence function to belong to a representor, see Joyce (2010), Mahtani (2016), Rinard (2017), and Levinstein (2019).

¹⁴ See Zynda (2000) and Steffánson (2017, 2018) for recent defenses of comparativism. An excellent survey by Konek (2019) discusses several versions of comparativism, as well as objections and responses to them.

¹⁵ Schoenfield (forthcoming) endorses this version of comparativism as applied to imprecise probabilities. See also Konek (2019) for discussion of extensions of comparativism to the imprecise case.

The Comparative View: A credence function c is a member of an agent's representor if and only if, for all propositions P , Q and R , if the agent is at least as confident in P given R as she is in Q given R , then $c(P|R) \geq c(Q|R)$, and if the agent is more confident in P given R than she is in Q given R , then $c(P|R) > c(Q|R)$.

The Comparative View has the nice result¹⁶ that given certain conditions on your comparative confidence ordering, there will be exactly one probability function in your representor; in other words, you will have precise credences. These conditions are that it satisfies certain intuitive axioms (like transitivity), that it is defined over “enough” propositions,¹⁷ and that it is complete, in the sense that for any two propositions, you are always either strictly more confident of one than the other, or equally confident in both. Imprecise agents, on this picture, have incomplete comparative confidence orderings, and the set of probability functions that represents them corresponds to the set of probability functions derived from all and only completions of the ordering.¹⁸

Given this understanding of imprecise doxastic states, we claim that a perfectly rational agent ought to guess as follows:¹⁹

¹⁶ For discussion of this and related results see Fishburn (1986), Luce (1968), Stefánsson (2017, p.578 and 2018, p.384-5), and Konek (2019).

¹⁷ This requirement can be satisfied in a number of ways (see Fishburn (1986) and Stefánsson (2017, 2018) for detailed discussion). What we'll be assuming (following results in Savage (1954), cited by Fishburn (1986, p.342), and Stefánsson (2018, p.385)) is that the ordering is “superfine”: whenever the agent is more confident in A than in the null proposition, the set of worlds over which the comparative confidence ordering is defined, can be partitioned into propositions $B_1 \dots B_m$, such that the agent is more confident in A than in B_i for all $B_i \in \{B_1 \dots B_m\}$. This requires an infinite algebra of propositions. (It's worth noting that uniqueness results *can* be obtained in a finite algebra as well, though the assumption required in that case, “Suppes Continuity”, as appealed to in Stefánsson (2017) is also a strong one). Although some might resist comparativism due to the fact that unique probabilistic representation is only obtained when the algebra is sufficiently rich, or structured in a certain way, our goal is not to defend comparativism here. Rather, we are *assuming* comparativism about probability, thereby taking on the board most of the assumptions needed for the relevant representation theorems, but relaxing the completeness assumption to account for imprecise probabilities. Our aim is to show how a comparativist picture of imprecise probability combined with an educated-guess understanding of accuracy can be used to respond to the dilation problem. Nonetheless, it's worth noting that requirements like the requirement that the algebra be super-fine are not unique to comparativism, but are present in much of the literature on subjective probability. For example, if instead of understanding subjective probabilities in terms of comparative confidence judgments, we understood subjective probabilities in terms of agents' betting dispositions, we'd need agents betting dispositions to be “superfine.” Additionally, Gibbard (2008)'s defense of Immodesty and Levinstein's (2019) elaboration of it depends on an agent's ability to consider “a continuum of bets”, making distinctions among bets whose payouts differ infinitesimally. Finally, this assumption it is present in Horowitz (2019)'s defenses of Probabilism and Immodesty. Thanks to anonymous referee for pressing us on this point.

¹⁸ Comparativism about imprecise probability also fits nicely with the fact that one way to motivate the claim that imprecise probabilities are needed to represent agents' belief states is by appeal to facts about agents' comparative confidence ordering: in particular, that in some cases an agent may be more confident in A^+ than in A , but no more in confident in A^+ than in B , and also no more confident in A than in B . No single probability function can represent this comparative confidence ordering.

¹⁹ Somebody who adopts what we'll call “the indeterminacy interpretation” of imprecise credences (Rinard (2015), Levinstein (2019)) may object to our licensing claims on the following grounds: On the indeterminacy interpretation, agents with imprecise credences are agents who have a precise credence function but it is indeterminate which precise credence function it is. The indeterminist may claim that it's indeterminate which guesses are licensed by an imprecise agent's state because it's indeterminate what the agent's (precise) credence function is, and different (precise) credence functions license different guesses. The spirit of our view is still

Simple Questions: A guess G_i in response to a forced choice question is forbidden if and only if there is some other guess G_j such that you are more confident in G_j than G_i . (In other words, every member of your representor assigns greater probability to G_j than G_i).

Suppositional Questions: A guess G_i in response to a forced choice question given a supposition S is forbidden if and only if there is some other guess G_j such that you're more confident in G_j given S than you are in G_i given S . (In other words: every member c of your representor is such that $c(G_j|S) > c(G_i|S)$).

We will argue only argue for **Simple Questions**, since the argument in favor of **Suppositional Questions** is exactly analogous. First, consider the right to left direction. Suppose you're more confident in G_j than G_i . In this case, it would always be irrational to guess G_i , given that there is a guess you regard as more likely to be true, namely G_j . Therefore, G_i is forbidden.

Second, consider the left to right direction. Suppose that there is no guess G_j such that you are more confident of G_j than of G_i . Then from your own point of view, you can't expect to do better than you would do by guessing G_i . So G_i should be permitted.

Let us give a few examples to show how this licensing norm works. Suppose your credence in P is the interval $[.4, .6]$. Then, if you are asked to choose between $\{P, \sim P\}$ you will be licensed to guess either one (since it is not true of you that $P > \sim P$ nor that $\sim P > P$; so, both guesses are permitted). Suppose further that your credence in Q is 0.2. Then, in response to the question $\{P, Q\}$, you will only be licensed to guess P ; in response to the question $\{\sim P, Q\}$ you will only be licensed to guess $\sim P$. Notice that there is no requirement that your guesses be consistent with one another – either in response to a single question (as we saw with $\{P, \sim P\}$ above) or in response to distinct questions.²⁰

With our licensing norm in place, we can prove the following general result: dilating one's doxastic state will result in strictly more guesses being licensed, and contracting one's doxastic state will result in strictly fewer guesses being licensed. More precisely, for any imprecise state M and any probability function p which is a member of M :

Contracting Preserves Guesses: Every guess licensed by p is also licensed by M .

compatible with (at least the letter of) the claim that an imprecise agent has a precise credence function but it's indeterminate which it is. If this is your view, then just think of our licensing norms as saying: If, *determinately*, you're more confident in P than Q , then you must guess P over Q . Otherwise, you're licensed to guess P and licensed to guess Q .

²⁰ An example suggested by an anonymous referee brings out this point. Suppose you are given a die that may or may not be loaded. You are completely clueless as to whether it is loaded or not, and as to how it is loaded, if it is. (Fill in your favorite story as to why you are so clueless.) So your credences about how it will land are maximally imprecise. In this situation, plausibly, you would be licensed to guess either way in response to all of the following questions: $\{1, \sim 1\}$, $\{2, \sim 2\}$, $\{3, \sim 3\}$, $\{4, \sim 4\}$, $\{5, \sim 5\}$, $\{6, \sim 6\}$. This means that, supposing you really had to answer all of those questions in a sequence, it would be fine for you to answer "1", then "2", then "3", and so on – even though you are well aware that there is no way all of your guesses, or even any two of them, could be true at the same time. But this strikes us as intuitively correct for an agent who really is maximally imprecise: after all, such an agent really does see the answer "1" as no better or worse than the answer " ~ 1 ", and the answer "2" as no better or worse than the answer " ~ 2 ", and so on. Recall that on this picture we are understanding agents as considering each question separately; the educated guess picture deliberately avoids evaluating *collections* of guesses, in response to different questions.

Dilating Adds Guesses: There are some guesses licensed by M that are not licensed by p .

This makes intuitive sense, given the way we are thinking about what imprecise credences are. Adding more credence functions to your representor means becoming less opinionated in terms of one's comparative confidence judgments. As one becomes more indifferent concerning how to compare propositions, one also becomes more indifferent concerning which guesses to make when forced to choose between them. As your doxastic state widens, from your own point of view, many guesses will seem to be no better or worse than one another. We provide a proof of these two claims in the appendix.

5. Dilation

We are now ready to argue for our first normative claim:

Don't Dilate: It is rationally impermissible to dilate.

Let us suppose that your current doxastic state is p , and you are thinking of moving to an imprecise state M , which has p as a member. From your current perspective, you'll think this is a bad move. More precisely, you will think that there is *some* question that you're more likely to get wrong if you move to M , and *no* question that you're more likely to get right.

Why? First recall that **Contraction Preserves Guesses**. This means that, for any question, every guess licensed by p is also licensed by M . But because **Dilation Adds Guesses**, there is some guess, in response to some question, that is licensed by M but *not* licensed by p . Call this not-licensed-by- p guess "G." Since G is not licensed by p , there is some alternative answer to the question, G^* , that *is* licensed by p , and which is such that $p(G^*) > p(G)$. So if I'm in p , I will think that the *only* thing that might go differently if I move to M , guessing-wise, is that I might make guesses that I regard as less likely to be true than the guesses I'd make if I stayed in p . Standard 'weak dominance' reasoning will therefore tell me that if I care about guessing truly, I should stay in p rather than move to M .

Here is a simple example. Suppose that you have a 0.4 credence that it will rain tomorrow, and you are asked to guess between: (A) It will rain tomorrow, and (B) A fair coin will land heads. Of course, your epistemic state only licenses guessing (B). You regard (A) as a less accurate guess. However, suppose you were to dilate from 0.4 to $[0.3, 0.7]$ on the rain. Then, your epistemic state *would* license a guess of (A). Because you currently consider (A) to be a strictly worse guess than (B), you shouldn't want to dilate if you want to be accurate.

6. Contraction

The second normative claim we will be arguing for is as follows:

Contraction is Fine: It is rationally permissible to contract.

Suppose you're in M and considering moving to p . If it's rationally impermissible on accuracy-based grounds to contract, you must in some sense expect to do better guessing with M than guessing with p . For this to be the case, there must be at least one question Q such that you think you're more likely to answer Q correctly using M than using p . But for you to regard it as more likely that you'll answer Q correctly using M , every credence function in M must be more confident that you'll answer Q correctly using M than using p . Since p is a member of M , it follows that p must be more confident that you'll answer Q correctly using M than using p . But given the considerations in the previous section, we know this isn't the case: There is no question which p thinks you're more likely to answer correctly using M than you are using p . After all, from the perspective of p , all M does, if anything, is *add* guesses which p thinks are less likely to be true than those licensed by p . So, you don't think it's more likely that you'll answer Q correctly using M . So it is permissible to contract.

As a corollary to **Contraction is Fine**, we get the result that, unlike precise credences, imprecise credences are *not* strictly self-recommending. In particular, by the argument just given, an imprecise doxastic state does not expect itself to be more accurate than any precise state contained in it.²¹

7. The Asymmetry Between Dilation and Contraction

Our results highlight an asymmetry between dilation and contraction: you should avoid dilating, but there's no reason not to contract. There are two interesting questions one might ask concerning the asymmetry:

Question 1: What explains the asymmetry between dilation and contraction on EGT? And why doesn't the asymmetry between dilation and contraction show up in EUT?

Question 2: Is this asymmetry acceptable, or unacceptable?

We'll address each in turn:

Question 1: Explaining the Asymmetry Between Dilating and Contracting

²¹ The asymmetry between dilating and contracting, as we've framed it, is an asymmetry between *becoming imprecise* and *becoming precise*. Before we move on to discuss the implications of this asymmetry, let us address a question you may be wondering about: is there a more general result we could prove, covering transitions *between* imprecise states as well? Do the arguments against becoming imprecise also support a ban on becoming *more* imprecise, or widening an already imprecise state?

These are interesting questions, which deserve more in-depth treatment. We will not pursue them in this paper, however, for a few reasons. Comparing different imprecise states to one another introduces all sorts of new questions and complications. Depending on how one sets things up, one can quickly run into problems: in particular, choosing whether to move from one state to another, when both of these states license multiple guesses based on the "opinion" of multiple probability functions, quickly raises tricky questions about judgment aggregation. We believe these issues will be better sorted out once we have settled the simpler cases.

The asymmetry between dilation and contraction in EGT shows up because EGT is based on the notions of *licensing* and *forbidding*. There are notable asymmetries between licensing and forbidding. To see this consider the following analogy:

Plans for Tonight: Picky and his friend Anything-Goes are going to spend an evening together. Anything-Goes will be happy with a wide range of activities: a wild dance-party, a trek in the jungle, a relaxing hang-out at the library, a nice meal, a not-so-nice meal – it all sounds great to Anything-Goes. Picky has very discriminating taste. He enjoys eating Cheerios, watching baseball, and not a whole lot else.

Picky and Anything-Goes should not have symmetrical attitudes towards one another. Anything-Goes should have no problem with Picky choosing the evening's activities, but Picky shouldn't be happy with Anything-Goes being in charge. After all, who knows what sort of nonsense Anything-Goes might choose?

The asymmetry between dilating and contracting is of the same sort. When you're in an imprecise state, you're like Anything-Goes: there will often be a number of guesses, in response to various questions, which you regard as no better or worse than one another. So you should have no problem with cutting down those options by moving to a state that licenses strictly fewer guesses. But when you're in a precise state, you *should* object to moving to one that is more permissive.

EUT, on the other hand, thinks of accuracy in terms of *distance* from the truth. It is unsurprising that an accuracy measure based on distance will not rule out either dilating or contracting. To see why consider the following silly analogy:

Dark Room: You're standing in a dark room. Your only goal is to be closer to the door, which is either on the wall to the left of you or to the right of you, but you don't know which. You're given the option of (painlessly) making your body wider, keeping your body at its current size, or making your body narrower.

Is there any reason to become wider or narrower? Is there any reason to stay the same size? If you only want to be closer to the door, it's hard to see why you'd prefer any of these options over any other. If you're wider, then no matter which side the door is on, some parts of you will be closer and others will be farther away. If you're narrower, some parts of you will be less far from the door than they would have been had you been wider, but other parts will be less close than they would have been. There's no clear sense in which you should expect to do better, relative to your goals, by either widening or narrowing. Similarly, there's no clear sense in which you should expect to do better relative to your accuracy goals, as understood by EUT, by dilating or contracting. This is an analogy, not an argument (for the arguments see our discussion earlier in the paper), but we hope it gives you a feel for why a distance-based notion of accuracy might struggle to motivate general prohibitions against dilating and contracting.

It makes sense, then, that one framework sees an asymmetry where the other does not. The explanation flows from the underlying notions which each framework uses to understand accuracy: licensing and forbidding on the one hand and distance on the other.

Question 2: Evaluating the Permissibility of Contracting

Let's turn now to assessing these results. As we've seen, *neither* EUT nor EGT can rule out contracting in general. Furthermore, EGT and the non-numerical approach to EUT we sketched earlier permit contraction to any member of the representor. Is this an undesirable result?

We will start with some potentially good news. Contraction of certain sorts may be helpful in alleviating some skeptical concerns. For example: Suppose you're a "superbaby", entering the world for the first time with no empirical information. Arguably, in such a state you are not justified in being confident that the world is any particular way: that the world is law-like, that certain predicates rather than others will project, that one's perception will be reliable, and so forth. Does this doom you to lifelong skepticism? An advocate of contraction might say no. For suppose the rational ur-priors are *imprecise*. Then you needn't remain skeptical. If contraction to any member of the representor is rationally permissible, it is perfectly fine to start off with an imprecise credence in the proposition that, say, perception is reliable, and then move to some precise high-ish credence that perception is reliable, even if you've been given no evidence in support of this conclusion! Once you start out with a bit of an anti-skeptical edge, then, depending, on how the rest of your experiences go, you might end up with quite ordinary common-sensical opinions.²² (Of course, it is also possible – and rational on this view – that the reverse might happen instead.)

But despite the potential anti-skeptical goodies, we imagine that many people will find the idea that contraction is rational to be highly unintuitive. One way to bring out the potentially problematic nature of contraction is to think about its potential practical consequences.²³ Consider:

Three Restaurants: You have a choice of going to restaurant A, B or C. You know that one of restaurants A and B is amazing, and one is terrible, but unfortunately you don't know which. You're certain however that restaurant C is very good.

Suppose you have no idea which of A or B is the amazing restaurant and that your attitude is represented by an imprecise credence $[a, b]$. Someone might worry as follows: "**Contraction is Fine** says you can move from $[a, b]$ to b , which might rationalize going to one of the risky restaurants. But that's a bad idea. Surely in this situation the only acceptable choice is to go to Restaurant C." This worry suggests that allowing contraction could be *practically* problematic.

Which judgments will seem appropriate in this case will depend on all three of the following: the values of a and b , the relative values of the restaurants, and your decision theory. In our view, however, in any case in which contraction from M to p will be permitted, *and* it makes sense to go to one of the risky restaurants given p , it will already have made sense to go to the risky restaurant before contraction, when you are in M . On the other hand, if post-contraction it is *not* worth the risk to go to restaurant A or B, then it was not worth the risk pre-contraction either. These verdicts are delivered by supervaluationist decision theories according to which an act is forbidden if and only if it is forbidden from the point of view of every member of your

²² For a similar way of motivating a non-skeptical response to certain higher order evidence cases, see Schoenfield (forthcoming).

²³ Thanks to Julia Staffel for pressing this point. The case below case is a less morbid version of the famous "miners puzzle" first published by Kolodny and McFarlane (2010).

representor. On this type of view, contraction will not permit any action that was not previously permitted.²⁴

Of course, you may well disagree with our judgments. You may think that in certain cases, an agent in M will not be rationally permitted to take risks like going to Restaurant A, whereas an agent in p will be so permitted. To deliver that verdict, you will need to adopt an *ambiguity-averse* decision theory. We think there are good independent reasons to reject ambiguity-averse decision theories, but we won't try to defend that claim here.²⁵ Instead, let's think about what the anti-contraction argument looks like, for someone who adopts an ambiguity-averse decision theory. We have argued that there are no *accuracy-based* reasons not to contract. The ambiguity-averse decision theorist should agree with us on this count. What she thinks is something like this: "if I contract, I won't be any less accurate than I am now. But I will be in a state that permits actions that, right now, look pretty bad to me. Better not do that!"

If we use an ambiguity-averse decision theory to avoid contraction, then we will have to separate the *accuracy* of your credences from their "guidance value" – roughly, how well they'll do at getting you what you want.²⁶ This isn't the worst result. Since neither EGT nor EUT says that accuracy considerations *require* contraction, it's not as though separating guidance value from accuracy will lead you to some sort of uncomfortable dilemma. But it is a nice feature of the accuracy frameworks on offer (both EUT and EGT combined with traditional or supervaluationist decision theory) that accuracy and guidance value go together – that what makes sense to believe for accuracy-based reasons also makes sense for guiding our actions.

Are there ways to rule out contraction, without separating accuracy and guidance value? One natural thought is that we could do this by amending the guessing framework, so that instead of only considering forced-choice scenarios, we also allowed the possibility of *abstention*. What if, in response to some questions, your doxastic state did not permit *any* guesses?²⁷ This might free up some space for prohibiting contraction, and for accepting ambiguity-averse verdicts in cases like **Three Restaurants**, without separating epistemic and practical concerns. We will turn to examining that possibility – and more generally, the prospects for *accuracy-based* arguments against contraction – in the final section.

²⁴ Making the case a bit more realistic might help. For one, notice that in any realistic case where you "have no idea" whether A or B is best, it's much more likely that you're representable by something like [0.4, 0.6] than by [0, 1]. (In [0, 1], not only do you lack a comparative confidence judgment between *A is best* and *B is best*. You also lack a comparative confidence judgment between *A is best* and $2+2=4$!) So suppose your credence in the proposition that A is best is [0.4, 0.6], and you contract to 0.6. Is the amazingness of the good restaurant worth the risk of the terrible restaurant, if you're only .6 confident of which is which? Maybe not. But in that case the problem we are worrying about does not arise. Maybe yes, it is worth it – but in that case, the amazing restaurant must be *really amazing*, the terrible restaurant *not that bad*, and so on; filling in the details would also make it more plausible that even in [0.4, 0.6], taking a gamble on Restaurant A is a sensible thing to do.

²⁵ But see Schoenfield (MS) and Al Najjar and Weinstein (2009).

²⁶ See Gibbard (2008) for a discussion of accuracy and guidance value, and the suggestion that they may come apart.

²⁷ Thanks to Alejandro Pérez Carballo for proposing this in conversation.

8. The Prospects For Accuracy-Based Arguments Against Contraction

First we'll explain why we don't think adding abstention as an option to the guessing framework will do the trick. Then we'll gesture towards some more general reasons why it's challenging to find accuracy-based arguments against contraction.

The problem with allowing abstention is that, if we do, we lose the result that precise credences are self-recommending. To put it another way: we lose the result that your guesses *elicit* your credences or your comparative confidence judgments, the way the Brier Score was designed to elicit weather forecaster's subjective probabilities. This is because, with abstention as an option, there will be cases where many different belief states, defined over the same algebra of propositions, deliver exactly the same guesses.

Here is an intuitive argument for why this is true. The main idea behind EGT is that your goal, in guessing, is to make true guesses and avoid making false ones. If we only consider forced choices, then when faced with two very improbable options, it makes sense for you to guess in favor of the option in which you have higher credence. So, your guesses will always elicit your comparative confidence judgments. But when we allow abstention – if the choice is *not* forced – then there will be many questions which don't elicit your comparative confidence judgments. For example, suppose you are asked to guess between:

S: It will snow in New York in September.

J: It will snow in New York in July.

If you have precise credences in S and J, it's safe to assume that your credence in S is higher than your credence in J. So in a forced choice between S and J, you will guess S. But if abstention is an option, and there's disvalue to guessing falsely, you may well abstain – after all, both S and J are almost certainly false. So the difference between your credence in S and your credence in J won't show up in your guesses: your guesses won't elicit your credences. This also means that you won't regard your own credences as the best for the purposes of guessing truly. This is because if the disvalue of guessing falsely is non-zero, there will be some number c such that, if you have credence less than c in P, you will not want to guess P, no matter what P is being compared with; it will always look better to abstain. So you won't have any accuracy-based reason to stick with your current credences, rather than adopting new ones, which swap out your credence in P for some other sub- c value.²⁸ This means that adding abstention as an option has the result that precise credences don't even recommend themselves over alternative precise credences.

²⁸ Are there other accuracy-based reasons not to make the switch – for example, because doing so would make you violate the probability calculus? Maybe some argument like this is available, especially if we can appeal to arguments from EUT *as well as* EGT. But in general it seems that adding the option to abstain will significantly weaken EGT. To take another example: if we move away from forced choices and allow abstention, Horowitz's (2019) argument for probabilism will not go through either. Part of that argument, for Boundedness, appeals to scenarios in which an agent is asked to choose between guessing the anti-tautology and guessing in favor of some other very improbable, contingent proposition like J, and relies on the thought that, in such situations, one should *never* guess in favor of the anti-tautology. But if the option to abstain is available, it seems that the most one could show here is that one's credence in the anti-tautology should be less than c . We would not be able to rule out the possibility that one's credence in the anti-tautology is *greater* than one's credence in J, which would violate Boundedness.

It seems then that adding abstention as an option to the guessing framework is a bad idea. Still, it's interesting to think about whether there is some plausible way of thinking about accuracy that yields the result that all probabilistic states – both precise and imprecise – are self-recommending.²⁹ While we have no general argument that this can't be done, we'll conclude with some admittedly hand-wavy suggestions as to why such a project might be difficult.

In general, belief states are self-recommending because they involve a commitment; they “take a stand”. In the case of binary belief, it's clear why belief in P recommends itself from the point of view of accuracy. To believe P just is to believe that P is true, which means that if you want a true belief, and you believe P, you'll think believing P is a good idea. We can also think of precise credal states as taking a stand. On the comparative picture, they're states that are committed to regarding certain propositions as more likely than others. On other pictures they may be states that are committed to regarding certain bets as preferable to others. Insofar as one is taking a stand, or making a commitment, that stand or commitment will naturally rule out competing commitments, just as an intention to ϕ rules out intentions to act in ways that are incompatible with ϕ -ing. But in general, while commitments rule out other commitments, a *lack* of commitment doesn't obviously rule out anything at all. (Compare: lack of intention regarding whether to ϕ doesn't rule out ϕ -ing.) The more we think of imprecise credal states as states in which certain commitments are *lacking*, as the comparative confidence picture does, the harder it will be to motivate the idea that imprecise states recommends themselves over their precise members.

So to rule out contraction, it might be better to think of imprecise credal states as expressing some form of commitment that rules out adopting their members. The practical motivation for imprecision we discussed in the previous section is an excellent example of this strategy: on ambiguity-averse decision theories, being in $[0,1]$ involves certain practical commitments that being at 0.5 (or any other precise state) does not have. That's why, on these theories, if you're at $[0,1]$ you will have practical reasons to stay there. An alternative approach might involve thinking of a state like $[0,1]$ as expressing some form of *evidential* commitment - perhaps a commitment to the effect that the evidence available is deficient in certain respects. But it's not clear what sort of commitment is expressed by $[0,1]$ that will rule out 0.5 on *accuracy* grounds.³⁰

²⁹ The two-tier lexicographic scoring rule in Seidenfeld et al. has this feature but it violates some of the constraints alluded to earlier – in particular, it violates the extensionality constraint (see Pettigrew (2016), Schoenfield (2017) and Berger and Das (2020) for a discussion of this constraint). To illustrate, it has the feature that a credence function which just assigns a 0.5 credence to each of Heads and Tails can be more accurate in Heads worlds than in Tails worlds.

³⁰ One way to go would be to think of the preference to remain imprecise as stemming primarily from certain features of an agent's epistemic *values* rather than just her doxastic commitments (or lack thereof). This is Konek's (forthcoming) strategy. He proposes rules that will motivate a preference to stay imprecise based on the idea that agents might prioritize having representors that *contain* accurate credence functions over having representors that *don't contain* inaccurate credence functions. But once the preference for staying imprecise comes from having *values* that favor imprecision in this way, the problem will be explaining why we should ever be precise. And indeed the versions of Konek's framework that favor imprecision don't allow for precision, and so precise states are no longer self-recommending. (This is why Konek's proposal doesn't satisfy the constraints appealed to in the results discussed earlier).

The considerations above are not of course a formal argument. Mathematics is full of complex structures and wiley tricks. We have no argument against the possibility of developing a new framework for thinking about accuracy which rules out both dilation and contraction,³¹ but philosophical reflection on the nature of imprecise states, we think, supports exactly the results delivered by the educated guess framework: the commitments involved in having precise credences result in a commitment to remain precise, but the lack of commitment involved in having imprecise credences doesn't ground a commitment to remain imprecise.³²

Appendix

Here are the two claims we will prove. For any imprecise state M , and any precise state p which is a member of M :

Contracting Preserves Guesses: Every guess licensed by p is also licensed by M .

Dilating Adds Guesses: There are some guesses licensed by M that are not licensed by p .

First we will prove that Contraction Preserves Guesses. Consider any arbitrary question Q , and suppose that p licenses guessing G in response to Q . Does M license guessing G ? Yes. For suppose M did *not* license guessing G . This would mean that every credence function in M must assign *lower* credence to G (conditional on any suppositions, if relevant) than to some alternative guess, G' . Since p is in M , this would mean that p must assign lower credence to G (conditional on any suppositions) than to G' . But then p would not license guessing G in response to Q . This contradicts our supposition that p licenses guessing G . Therefore, if p licenses guessing G in response to Q , M does too.

On to **Dilating Adds Guesses**. The form of our argument is as follows:

- (1) An agent in p has a comparative confidence judgment that an agent in M lacks.
- (2) If an agent in p has a comparative confidence judgment that an agent in M lacks, then there is some guess that is licensed by M and not licensed by p .

Let's start with Premise 1, which we'll defend using the comparativist picture of imprecise credences that we favor. Since M is an imprecise state, there is at least one pair of propositions in the algebra, such that M lacks a comparative confidence judgment between them. Since p is precise, p has a comparative confidence judgment for any pair of propositions in the algebra. It follows that p has a comparative confidence judgment that M lacks. Suppose that the

³¹ Though as the results discussed in earlier sections show, we do believe that this can't be done in a plausible way by assigning belief states accuracy *scores* as is done in EUT.

³² For helpful comments, suggestions, and questions, we are grateful to Zach Barnett, Jason Konek, Chris Meacham, Alejandro Pérez Carballo, Richard Pettigrew, Susanna Siegel, Julia Staffel, and audiences at Ohio State University, the National University of Singapore, the Harvard Bounded Rationality Workshop, The Vagueness, Parity and Non-Conventional Comparative Relations Conference at the Institute for Future Studies and The International Conference for Philosophy of Science and Formal Methods for Philosophy at the University of Gdansk. We would also like to thank two anonymous referees, for their helpful and extensive comments.

comparative confidence judgment that p has and M lacks is between A and B . Then, the comparative confidence judgment will be one of the following: $A = B$, $A > B$, or $A < B$.

On to Premise 2: We will show that whatever that judgment is, M will license some guesses which p does not.

Let's suppose first that p is strictly *more* confident in A than B , but M is not. Then p will forbid guessing B in a choice between A and B , but M will license guessing B (as well as A) when faced with such a choice. So if p is more confident of A than B , M will license a guess that p does not license. (A parallel argument goes through if p is strictly less confident of A than B , but M is not.)

Now suppose that p is equally confident in A and B , but M is not. In this case, p and M will both license the same guesses in response to the question $\{A, B\}$. However, we will show that in this scenario, there is a *different* question, in response to which M will license strictly more guesses than p .

Since M is not equally confident in A and B , there is some credence function $p^* \in M$ doesn't assign equal credence to A and to B . Let $p(A) = p(B) = r$. Now let p^* be some credence function in M that is not p . Without loss of generality assume $p^*(B) > p^*(A)$.

Case 1: $p^(A) \geq r$*

Since $p^*(B) > p^*(A)$, it follows that $p^*(B) > r$. Let $r+\epsilon$ be a number such that $p^*(B) > r+\epsilon > p^*(A)$.

Now consider the question: $\{B, \text{an } r+\epsilon\text{-weighted-coin lands Heads}\}$.

p will be more confident that the $r+\epsilon$ -weighted-coin lands Heads.

p^* will be more confident that B .

So there is a guess licensed by p^* (namely B) that is not licensed by p . Since any guess licensed by p^* is licensed by M , we have a guess that is licensed by M but not licensed by p .

Case 2: $p^(A) < r$*

Let $r-\epsilon$ be a number such that $p^*(A) < r-\epsilon < r$.

Now consider the question: $\{A, \text{an } r-\epsilon\text{-weighted-coin lands Heads}\}$.

p will be more confident that A .

p^* will be more confident that the $r-\epsilon$ -weighted-coin lands Heads.

So there is a guess licensed by p^* (namely an $r-\epsilon$ -weighted-coin lands Heads) that is not licensed by p . Since any guess licensed by p^* is licensed by M , we have a guess that is licensed by M but not licensed by p .

Since either $p^*(A) \geq r$ or $r > p^*(A)$, it follows that there is some guess that is licensed by M but not licensed by p .

It follows that **Dilating Adds Guesses**: for any imprecise state M and precise state p that is in M 's representor, there are some guesses licensed by M that are not licensed by p .

References

- Berger, D. & Das, N. (2020). Accuracy and Credal Imprecision. *Noûs*. Advanced online publication. doi: 10.1111/nous.12274
- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review* 122: 337–93.
- Builes, D. & Spencer, J. (MS) How (Not) to Justify Chance-Credence Norms.
- Elga, A. (2010). Subjective Probabilities Should be Sharp. *Philosophers Imprint* 10, no. 5: 1–11.
- Fishburn, P. (1986). The Axioms of Subjective Probability. *Statistical Science* 1: 335–358.
- Gibbard, A. (2008). Rational Credence and the Value of Truth. In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, 143–164. Oxford University Press.
- Horowitz, S. (2019). Accuracy and Educated Guesses. In T. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology Volume 6*, 85–113. Oxford University Press.
- Jeffrey, R. (1983). Bayesianism With A Human Face. In J. Earman (Ed.), *Testing Scientific Theories*, 133–156. University of Minnesota Press.
- Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives* 19: 153–178.
- Joyce, J. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, 263–297. Synthese.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24: 281–323.
- Kaplan, M. (2010). In Defense of Modest Probabilism. *Synthese* 176: 41–55.
- Kolodny, N. & McFarlane, J. 2010. Ifs and Oughts. *The Journal of Philosophy* 107: 115-143.
- Konek, J. (2019). Comparative Confidence. In *The Open Handbook of Formal Epistemology*, edited by Richard Pettigrew and Jonathan Weisberg, 267-348. PhilPapers Foundation.
- Konek, J. (Forthcoming). Epistemic Conservativity and Imprecise Credence. *Philosophy and Phenomenological Research*.

- Levi, I. (1985). Imprecision and Indeterminacy in Probability Judgment. *Philosophy of Science* 52: 390–409.
- Levinstein, B. (2017). A Pragmatist’s Guide to Epistemic Utility. *Philosophy of Science* 84: 613–638
- Levinstein, B. (2019). Imprecise Epistemic Values and Imprecise Credences. *Australasian Journal of Philosophy* 97: 741–760.
- Luce, R. D. (1968). On the Numerical Representation of Qualitative Conditional Probability. *The Annals of Mathematical Statistics* 39: 481–491.
- Maher, P. (2002). Joyce’s Argument for Probabilism. *Philosophy of Science* 69: 73–81.
- Mahtani, A. (2016). Imprecise Probabilities and Unstable Betting Behavior. *Noûs* 52: 69–87.
- Mayo-Wilson, C. & Wheeler, G. 2016. Scoring Imprecise Credences: A Mildly Immodest Proposal. *Philosophy and Phenomenological Research* 93: 55–78.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press UK.
- Rinard, S. (2015). A Decision Theory for Imprecise Credences. *Philosophers Imprint* 15, no. 7: 1–16.
- Rinard, S. (2017). Imprecise Probability and Higher Order Vagueness. *Res Philosophica* 94: 257–273.
- Savage, L. (1954). *The Foundations of Statistics*. Wiley.
- Schoenfield, M. (2017). The Accuracy and Rationality of Imprecise Credences. *Noûs* 51: 667–685.
- Schoenfield, M. (Forthcoming). Meditations on Beliefs Formed Arbitrarily. In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology Volume 7*. Oxford University Press.
- Schoenfield, M. (MS). Can Imprecise Probabilities Be Practically Motivated?
- Seidenfeld, T., Schervish, M.L, & Kadane, J.B. (2012). Forecasting with Imprecise Probabilities. *International Journal of Approximate Reasoning* 53: 1248–61.

Stefánsson, H. (2017). What Is ‘Real’ in Probabilism? *Australasian Journal of Philosophy* 95: 573–587.

Stefánsson, H. (2018). On the Ratio Challenge for Comparativism. *Australasian Journal of Philosophy* 96: 380–390.

Sturgeon, S. (2008). Reason and the Grain of Belief. *Noûs* 4: 139–165.

White, R. (2009). Evidential Symmetry and Mushy Credence. In T. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, 161–186. Oxford University Press.