An Epistemic Argument for Stare Decisis

Bob Beddor

1 The Project

A central legal principle is *stare decisis*, the doctrine that judicial decisions should be constrained by precedent. It is hard to overstate the importance of *stare decisis* in the legal system. Flip through any court opinion and you will find it filled with discussions of precedents bearing on the case at hand. In the vast majority of cases, these discussions are not just window dressing. Rather, the court's assessment of how precedent bears on the facts at hand is one of the main factors determining the resolution of the case.¹

Of course, precedents are not immutable; they sometimes get overturned. There is an ongoing debate about the conditions under which courts should break with precedent— a question that has gained renewed urgency in the wake of the US Supreme Court's decision in *Dobbs v. Jackson Women's Health Center*, which upended nearly 50 years of precedent by overruling *Roe v. Wade*. However, virtually all parties agree that precedent should carry *some* weight.

But why? Why should precedent carry *any* weight? What normative justification can we provide for this principle?

Virtually all discussions of this question have focused on the practical benefits of *stare decisis*. Perhaps the most common rationale emphasizes the benefits of predictability. In a system where future decisions are constrained by precedent, people can make predictions about the law on the basis of past rulings. They can then rely on these predictions in their behavior—for example, in deciding whether to pursue a lawsuit, or to engage in behavior that might incur one.² Another prominent argument focuses on judicial efficiency. By deferring to precedents, judges can avail themselves of the hard work of other judges, rather than deciding each case from scratch.³

These pragmatic arguments capture important benefits of *stare decisis*, to be sure. But they leave one wanting more. Put yourself in the shoes of a judge tasked with ruling

¹See e.g., the 9th Circuit Court of Appeal's statement, "As every first-year law student knows, the doctrine of *stare decisis* is often the determining factor in deciding cases brought before any court." *Or. Natural Desert Ass'n v. U.S. Forest Serv.*, 550 F.3d 778, 785 (9th Cir. 2008).

²See e.g., Goodhart 1934; Wasserstrom 1961; Greenawalt 1977. This rationale also plays a starring role in economic analyses of law; see e.g. Kornhauser 1989. Considerations of reliance also feature prominently in courts' own discussions of the costs of overturning precedent; see e.g., *Planned Parenthood of Southeastern Pa. v. Casey*, 505 U.S. 833; *Dobbs v. Jackson Women's Health Organization* 142 S.Ct. 2228. For further discussion of this rationale, see Duxbury 2008: chp.3; Waldron 2012.

³See e.g., Stevens 1983; Schauer 1987; Duxbury 2008: chp.3; Burton 2014.

on some point of law. You are sure that there is a precedent p on this point. You are faced with the decision of deferring to p or breaking with precedent. Now, you can run through all the pragmatic rationales for deferring to p until you are blue in the face. A further question remains: what is *correct*, as a matter of law? Is the mere fact that p is precedent any reason to think that p is in fact true?

This paper argues that—under certain conditions—the answer is *yes*. I show how philosophers of law and legal scholars can draw on work in epistemology to argue that *stare decisis* maximizes the expected accuracy of a court's decisions. That is, a rational court should expect that overturning precedents, absent some positive reason to do so, will result in less accurate decisions. The result is a purely epistemic argument for deferring to precedent.

Before plunging in, let me pause to make a remark about the larger implications of this project. Recent years have seen a flourishing of work at the intersection of epistemology and law. For the most part, this work has focused on specific legal issues where epistemology is obviously relevant—for example, the use of statistical evidence in criminal and tort law, and the epistemology of eyewitness testimony.⁴ These issues, while doubtless important, only arise in a fraction of cases. Given this, one might get the impression that epistemology is only relevant to niche legal topics. My argument shows that this impression is mistaken. Epistemology has an important role to play in justifying *stare decisis*. And *stare decisis* permeates virtually every area of the law; it is something that no lawyer, judge, or legal scholar can afford to ignore.

2 Analogous Principles in Epistemology

It's common to distinguish between 'horizontal' and 'vertical' *stare decisis*. Horizontal *stare decisis* refers to the idea that a court should not revise its own precedents, absent a positive reason to do so. Vertical *stare decisis* refers to the idea that lower courts should defer to precedents set by higher courts, as when, say, a U.S. district court is constrained by the prior rulings of the U.S. Supreme Court.

My epistemic argument for *stare decisis* starts by showing that both forms the doctrine bear striking similarities to principles defended in epistemology.

2.1 Horizontal Stare Decisis and Epistemic Conservativism

It's generally thought that horizontal *stare decisis* carries less weight than vertical *stare decisis*. In legal lingo, the former is 'persuasive' rather than 'binding'; it can be overridden, at least in some circumstances. For present purposes, we need not weigh into the thorny debate about the conditions under which it can be overridden (though we will return to this question later). For now, we can operate with a fairly minimal construal of the doctrine:

⁴See e.g., Tribe 1971; Enoch et al. 2012; Buchak 2014; Smith 2018; Ross 2019; Bollinger forthcoming, among many others.

HORIZONTAL *STARE DECISIS* A court should not revise its own precedents, absent a positive reason to do so.

Horizontal *stare decisis* resembles a principle in epistemology sometimes goes by the name of 'epistemic conservativism'. According to this principle, a rational agent should retain their current beliefs and credences unless they receive some positive reason (usually in the form of new evidence) to rethink them. That is:

EPISTEMIC CONSERVATIVISM A rational agent should not revise their doxastic states absent a positive reason to do so.⁵

Epistemic Conservativism seems very plausible. Suppose that at noon Tanya believes that an asteroid caused the extinction of the dinosaurs. Between noon and 1pm, Tanya gains no new paleontological evidence. However, when the clock strikes one, she abandons her belief: she now suspends judgment on the causes of dinosaurian demise. There seems to be something irrational about her doxastic flipflopping.

Some might think this is too quick. What if Tanya had no good reason for initially believing that an asteroid caused the extinction of the dinosaurs? If so, then her abandonment of this belief seems like good epistemic housekeeping: she is tossing out the doxastic dross (cf. Foley 1983; Christensen 1994). However, I think on closer examination this apparent counterexample loses force. We need to ask: why did Tanya abandon her original belief? If the answer is that she realized she held it for no good reason, then this realization itself qualifies as a positive reason for revising her belief, in which case her belief abandonment conforms to Epistemic Conservativism. By contrast, if her belief abandonment was not prompted by any such realization, then her doxastic seesawing looks much worse. In that case she will have sinned twice over: once by irrationally forming a belief, once again by irrationally abandoning it.⁶

So Epistemic Conservatism has considerable plausibility, and it has ample resources for avoiding immediate counterexamples. It also fits naturally with Bayesian epistemology. According to Bayesians, a rational agent should update their credences by condi-

⁵For explicit defenses of this idea, see Chisholm 1980; McCain 2008. Even when not explicitly endorsed, versions Epistemic Conservativism can be discerned in many different authors. According to Quinean holism, we seek to accommodate new experiences in ways that conserve as many of our old beliefs as possible (Quine 1980). According to the methodology of reflective equilibium, we should seek to work through our considered beliefs to achieve coherence and stability (Rawls 1971). And according to dogmatists and phenomenal conservatives, the mere fact that some seems to you to be the case is—at least in certain domains—a *prima facie* reason for believing it (Pryor 2000; Huemer 2007).

⁶Another response to this objection would be to recast Epistemic Conservatism as a wide scope norm:

WIDE SCOPE INTERPRETATION If one is a rational agent, one ought not both: (have some doxastic state \mathcal{D} & revise \mathcal{D} without some reason to do so).

Thus construed, Epistemic Conservatism is silent on whether the fault lies in Tanya's original formation of her belief or in her abandonment thereof.

tionalizing on their evidence.⁷ From this diachronic norm it trivially follows that if a rational agent gets no new evidence between t_1 and t_2 , they should not revise their credences between t_1 and t_2 . It also follows that a rational agent should not revise their *beliefs* between t_1 and t_2 if we adopt a 'Lockean' view of rational belief:

LOCKEANISM An epistemically rational agent believes p iff they assign p a sufficiently high credence.⁸

To my knowledge, no one has previously remarked on the connection between horizontal *stare decisis* and Epistemic Conservativism. The former principle has been mostly been debated by legal scholars and philosophers of law; the latter by epistemologists. But the two principles are strikingly similar. Both are *inertial* principles: they enjoin us to stick with our previous commitments (precedents in one case, beliefs in the other), absent some special reason to rock the boat. This close parallel suggests an intriguing possibility: perhaps we could derive the former from the latter.

In the paper, I show how this can be done. I spell out the details in §3, but the key idea is to think of judicial decisions as expressing a court's beliefs on the relevant legal questions. Consequently, horizontal *stare decisis* is just a special case of Epistemic Conservativism.

2.2 Vertical Stare Decisis and Expert Deference Principles

Whereas horizontal *stare decisis* is considered 'persuasive', vertical *stare decisis* is considered 'binding': courts are required to strictly follow the precedents set by higher courts in the same jurisdiction. For our purposes, we can understand vertical *stare decisis* as follows:

VERTICAL STARE DECISIS A court should defer to precedents set by higher courts.

This idea resembles what are known as 'expert deference' principles:

EXPERT DEFERENCE Upon learning that an expert believes p, a rational agent should defer to that expert and also believe p.⁹

⁷That is, let c be an agent's prior credence function, c' their credence at some later time, and E be a proposition representing their total evidence at this later time. Then conditionalization requires that $c'(\cdot) = c(\cdot | E)$.

⁸For defenses of Lockeanism, see Foley 1993; Christensen 2005; Sturgeon 2008; Leitgeb 2014; Easwaran 2016; Dorst 2019. As I am understanding Lockeanism, the view is a normative claim; it is noncommittal on the metaphysical question of whether belief just is sufficiently high credence. (Though obviously those who endorse the metaphysical identity claim should also accept the normative claim.) My construal is also neutral on issues of the order of explanation—that is, whether the rationality of full belief is explained in terms of the rationality of credence or *vice versa*. For discussion, see Easwaran 2016; Dorst 2019.

⁹There are a number of different versions of expert deference principles floating around in the literature. One version is the Reflection Principle (van Fraassen 1984), which says that you should treat you future self as an expert; another is Principal Principle (Lewis 1980), which says that you should treat the objective chance function as an expert. For recent discussion and defense of deference principles, see Dorst 2020; Dorst et al. 2021; Levinstein forthcoming.

Expert Deference is very plausible. Suppose you have no opinion about whether it will rain. But your friend Ella is a celebrated meteorologist, with access to the latest forecasting models. Ella tells you that she believes it will rain. Intuitively, you should defer to her.

Some might worry that Expert Deference is too strong. Suppose I have evidence that the expert lacks. For example, I know that Ella believes that it is not raining now. But Ella has been in a windowless room all day, and I am outside, caught in a surprise downpour. Surely, I should not defer to Ella's beliefs on the matter. However, we can avoid this worry by both (i) restricting expertise to particular topics, (ii) more carefully defining what it means to be an expert about some topic. For me to properly regard Ella as an expert about some topic, I must expect Ella to be more likely to be right than me about that topic.¹⁰ In this scenario: I properly regard Ella as an expert about the weather in general, but I should not regard her as an expert about whether it is raining right here, right now.

Thus clarified, Expert Deference holds considerable appeal. And it bears an obvious resemblance to vertical *stare decisis*. This suggests an intriguing possibility. Perhaps we can derive the legal doctrine from the corresponding epistemological principle.

In the next section, I make good on this possibility. To preview: the key idea is that lower courts should regard higher courts as *legal experts*. Consequently, vertical *stare decisis* is itself a type of expert deference principle, instructing a court to conform its (judicial) opinions to the (judicial) opinions of the legal experts. The upshot is that we can derive both forms of *stare decisis* from independently plausible principles in epistemology.

3 Deriving Stare Decisis

3.1 Judicial Rulings as Expressions of Beliefs

Before we can derive *stare decisis* from the corresponding epistemological principles, we need one more ingredient. Specifically, we need to say something about the relation between what court *rules* and what the court *believes*. Here is, I think, a very plausible thought:

RULINGS EXPRESS BELIEFS If a court issues a ruling on some point of law p, that ruling expresses a court's belief that p.

We can support this thought in two steps. The first step is to note that judicial rulings are *assertions*. As evidence for this, note that we regularly report judicial opinions by saying that the court "said" such-and-such—a locution that we typically reserve to report assertions. Some examples "from the wild":

¹⁰We'll see how to spell this out with more precision in §5.

- (1) The Supreme Court says colleges and universities can no longer take race into consideration as a specific basis for granting admission...¹¹
- (2) The Supreme Court said in *Cooper v. Aaron* (1958) that Article VI requires state government officials to enforce the Constitution as interpreted by the Supreme Court.¹²

The second step is to appeal to the widely endorsed idea that assertion is a speech act that expresses a speaker's belief. Rulings Express Beliefs is just a special case of this idea, where the relevant assertion is a ruling.

What does it mean for a ruling to *express* a belief? The short answer: whatever it means for an assertion to express a belief. Different speech act theorists give different accounts of the "expression" relation. But one natural approach is to unpack this in normative terms. According to this approach, to say an assertion expresses a belief is to say that assertion is subject to a sincerity norm:

SINCERITY NORM An agent should only assert p if they believe p.

If we accept this account of "expression", then Rulings Express Beliefs amounts to the view that judicial rulings are also governed by a sincerity norm: a court should only issue some ruling p if they believe p. For my purposes, I'll help myself to this account of the expression relation going forward, though much of what I say can be adapted to other ways of understanding this relation.

Some might grant that judicial rulings are assertions, but insist that they are a special type of assertion, one where the normal connection with belief is severed. Two lines of argument undercut this response. First, we use "belief"-y language to report judicial decisions. Courts describe their decisions as "judgments", which are announced in "opinions." Similarly, legal scholars, government agencies, news outlets, and courts themselves regularly make claims about what various courts "believe" on the basis of these opinions. Examples are legion, but here are a few:

- (3) This would seem to suggest that while the Supreme Court believes [the *Purcell* rule] applies to lower courts, it is not a constraint on its own rulings.¹³
- (4) [B]ecause this case was remanded for reconsideration in light of Alabama v. Smith, this Court must assume that the Supreme Court believes that Smith affects the instant case in some way...¹⁴
- (5) This Court believes that the First Circuit decision represents the sounder view.¹⁵

¹¹Source: CNN Politics

¹²Source: https://www.scu.edu/ethics-spotlight/the-ethics-of-guns/the-right-of-state-governmentsto-defy-the-supreme-court/

¹³Source: *Politico*, "On the Supreme Court's "Breathtakingly Radical" New Approach to Election Law."

¹⁴Source: Turner v. State of Tenn., 726 F. Supp. 1113: 1115.

¹⁵Source: Waters v. Plyborn, 93 F. Supp. 651 (E.D. Tenn. 1950)

What entitles us to conclude that a court believes p from the fact that they said p? Rulings Express Beliefs gives us an answer. We assume courts are abiding by the Sincerity Norm, and would only issue a ruling if they believed it.

A second line of argument is more indirect. A powerful source of evidence that assertions express beliefs comes from Moore's paradox (Moore 1962). Moore famously observed that there seems to be something incoherent about saying:

(6) ?? It's raining, but I don't believe it's raining.

A standard explanation for this infelicity comes from the idea that assertions express beliefs. Cashed out in terms of the Sincerity Norm: a speaker ought only assert (6) if they believe it. Since believing a conjunction requires believing both conjuncts, a speaker ought only assert (6) if they believe it's raining. But then then the second conjunct (*I don't believe it's raining*) is guaranteed to be false. So no one could truly utter a sentence such as (6) while abiding by the Sincerity Norm.¹⁶

This argument extends to the assertions made by courts. Take any holding you like. Consider, for example, *Loving v. Virginia*, where the US Supreme Court held that state laws banning interracial marriage violates the Due Process and Equal Protection Clauses of the 14th Amendment. Chief Justice Warren's opinion concludes with the soaring pronouncement that "[u]nder our Constitution, the freedom to marry, or not marry, a person of another race resides with the individual, and cannot be infringed by the State." Now, imagine if this pronouncement had been conjoined with an admission that the Warren court did not believe the relevant proposition:

(7) ?? The freedom to marry cannot be infringed by the State, but I [/we] don't believe this.

This sounds terrible—just as bad as any old Moorean assertion. So the Moorean argument for thinking that assertions express beliefs also supports thinking that judicial rulings express beliefs.

Summing up: there is considerable support for the idea that judicial rulings express beliefs of the court. We now have all the tools we need to derive both forms of *stare decisis* from the corresponding epistemological principles.

3.2 Vindicating Stare Decisis

Start with Horizontal *Stare Decisis*. Consider a court C which is tasked with ruling on some point of law. Imagine that there are no precedents constraining C's ruling on this point. After deliberating, C issues a ruling on the matter at time t_1 : C holds p. In doing

¹⁶Some philosophers have used Moorean assertions involving knowledge (e.g., "It's raining, but I don't know it's raining") to motivate a stronger norm, according to which a speaker should only assert p if they know p (e.g., Unger 1975; Williamson 2000). All of my arguments here are compatible with this stronger norm, but do not require it.

so, C expresses a belief that p. Suppose C is sincere: C does believe p at t_1 . Now consider some time t_2 when C is confronted with a new case, which raises the same question of law. By Epistemic Conservativism, it follows that C should not revise their belief in pbetween t_1 and t_2 absent some positive reason to do so. So if at $t_2 C$ were to overturn their earlier decision by ruling $\neg p$ (or, more generally, something inconsistent with p), absent a positive reason to do so, they would be violating either Epistemic Conservatism or the Sincerity Norm.

An analogous argument can be used to undergird vertical *stare decisis*, at least in certain circumstances. Suppose C is a lower court, and learns that a higher court issues some ruling p. Given the assumption that the higher court is sincere, C can conclude that the higher court believes p. In many circumstances, lower courts should regard higher courts as experts on matters of law. (There may be some situations were this does not apply; I'll return to this later.) Assuming C regards the higher court as an expert, Expert Deference says that C should defer to this higher court, and also believe p. And so C should not issue rulings that are inconsistent with p, on pain of insincerity.

3.3 Courts as Group Agents

Let me pause to address an immediate concern that may be weighing on some readers. Many courts are composed of multiple judges. One might worry that talking about what a court "believes" is all well and good when a court consists of a single judge, but makes no sense when we are dealing with a multi-member court.

However, by now many philosophers accept that an organized collective of individuals can qualify as intentional agent (French 1984; Rovane 1997; List and Pettit 2011; List 2018; Tollefsen 2015). Standard examples of such "group agents" include firms, corporations, scientific advisory boards, governments, and—most importantly for our purposes courts. These groups engage in sophisticated, purposeful behavior, in much the same the way that individual agents do.

Group agents can hold beliefs. For example, we might truthfully say:

(8) The Intergovernmental Panel on Climate Change believes that global warming is a serious threat.

The same applies to courts. Indeed, all of our earlier examples of ascriptions of judicial beliefs ((3)-(5)) were to multi-member courts. So there is no principled problem talking about what a multi-member court "believes".

Attending to the properties of group agents also puts to rest a related concern for my argument in §3.2. The concern is this: courts are not just constrained by their own precedents and those of higher courts. They are also constrained by precedents set by their predecessors in the same court. Some might worry that my argument for horizontal *stare decisis* fails to capture this constraint.

In response, it will help to draw attention to another notable feature of group agents. They can endure over time, despite changes in their composition. For example, the U.S. Supreme Court continues to exist even when one—or all—of its members is replaced. This view comports nicely with the way courts conceive of their own persistence conditions, as reflected in courts' penchant for using first person plural pronoun "we" to describe decisions that were issued by the same court decades or even centuries earlier. This provides a simple response to the worry under consideration. When a judge—or group of judges—defers to the precedents set by their predecessors on the same bench, the court is simply deferring to its earlier beliefs.

4 Foundations in Epistemic Decision Theory

4.1 A Concern About the Argument Thus Far

I have argued that *stare decisis* can be derived from independently plausible principles in epistemology: Epistemic Conservativism and Expert Deference. Some might worry that this just pushes our starting question back a step. Recall our starting question: You're a judge, tasked with ruling on some point of law. Does the fact that p is precedent give you any reason to think that p is more likely to be true? My answer so far boils down to the following: either p is one of your previously held beliefs, or it is a belief held by an expert. Either way, we have a plausible epistemic principle that says you ought to defer to p—Epistemic Conservatism in the former case, Expert Deference in the latter. But now our initial question resurfaces. Do we have any reason to think that abiding by these epistemic principles make your beliefs more likely to be accurate?

We can give this worry more teeth. When Epistemic Conservatism and Expert Deference were first put forward in the literature, much of the early discussion concerned their putative pragmatic benefits. For example, Sklar 1975 and Lycan 1988 suggest that Epistemic Conservatism might be motivated by considerations of cognitive efficiency. Revising our beliefs takes effort, so sticking with the doxastic *status quo* will reduce psychological energy expenditure. This style of argument bears a striking resemblance to the pragmatic rationales traditionally given for *stare decisis*. As noted in §1, a number of legal scholars have likewise justified *stare decisis* on the basis of efficiency considerations.

Similarly, one of the most influential early arguments for a version of Expert Deference was van Fraassen's dutch book argument for Reflection (van Fraassen 1984), which showed that if an agent does not plan to defer to the opinions of their better-informed future self, they will be vulnerable to a "Dutch strategy"—that is, an algorithm for placing bets that guarantees a profit to one side. While this differs in detail from the sort of practical considerations used to justify *stare decisis*, it is similarly pragmatic in nature. So an obvious worry looms. If the ultimate justifications for our epistemic principles are pragmatic, then our goal of providing a purely epistemic rationale for *stare decisis* still hovers out of reach.¹⁷

¹⁷See Christensen 1994 for an objection to pragmatic rationales for Epistemic Conservativism in this vein. Christensen develops this objection with an analogy to Pascal's wager. Granted, Pascal's wager may show that belief in God is practically beneficial. But it does nothing to show that such a belief has epistemic

But there's good news. Recent work in epistemic decision theory offers a promising framework for developing purely epistemic rationales for epistemic norms.¹⁸ This section provides an overview of epistemic decision theory and shows how it can be used to underwrite Epistemic Conservatism and Expert Deference. By combining this framework with the argument in §3, we get a purely accuracy-based vindication *stare decisis*.

4.2 Epistemic Decision Theory and Epistemic Conservatism

Epistemic decision theory starts with the idea we want to *get things right*: we care about having accurate beliefs (and degrees thereof). Insofar as we are epistemically rational, then, we will adopt whatever doxastic states maximize expected accuracy:

Acc-Max An epistemically rational agent should maximize the expected accuracy of their doxastic states.

How should measure the accuracy of a doxastic state? This is controversial, particularly when the relevant doxastic state is a credence function. In the literature, a variety of rules for 'scoring' the accuracy of a credence function have been put forward. For our purposes, we do not need to pick any particular scoring rule. Rather than singling out any single correct scoring rule, we can discuss constraints on any admissible scoring rule—that is, constraints that any rational way of scoring accuracy will satisfy. In the literature, one of the most widely accepted constraints is:

STRICT PROPRIETY The expected accuracy of a credence function c, when calculated using c, is strictly greater than the expected accuracy of any alternative credence function c'.¹⁹

Virtually all of the main measures of accuracy are strictly proper. A common rationale for Strict Propriety runs as follows. Suppose, by the lights of your current credence function c, some other credence function c' had higher expected accuracy than yours.

$$EA^{c'}(c) = \sum_{w \in W} c'(w)A(c,w).$$
(1)

A strictly proper accuracy measure is one where $EA^{c'}(c)$ is maximized when c' = c.

merit—that it is likely to be true, or supported by the evidence. Similarly, Christensen objects, we might grant that abiding by Epistemic Conservatism is practically beneficial. But do we have any reason to think that it also carries epistemic benefits?

¹⁸Epistemic decision theory is a movement in formal epistemology that adapts the tools of rational choice theory to justify epistemic norms. Notably, one major impetus for developing epistemic decision theory was a dissatisfaction with pragmatic rationales for various epistemological constraints—e.g., dutch book arguments for probabilism. For a sampling of important contributions to the program, see Joyce 1998, 2009; Greaves and Wallace 2006; Moss 2011; Pettigrew 2016; Schoenfield 2016; among many others.

¹⁹Formally: let an accuracy score A be a function from a credence function c and a world w to the real numbers. The expected accuracy of a credence function c, as calculated by some credence function c', is given by:

Then your credences would be self-undermining. Given Acc-Max, epistemic rationality would require you to abandon c in favor of c'.²⁰

Combining Acc-Max with Strict Propriety yields a simple justification for Epistemic Conservativism. If you are rational, then by your lights your credence function has greater expected accuracy than any alternative credence function (by Strict Propriety). So you should expect that, in the absence of new evidence, switching to some alternative credence will carry a loss in accuracy. By Acc-Max, you should stick with you current credence function.

I should note that from the perspective of epistemic decision theory, this is a pretty uncontroversial result. A number of authors in the literature have used the resources of epistemic decision theory to argue for much stronger diachronic norms. For example, Greaves and Wallace 2006 offer a proof that, given Strict Propriety, it maximizes expected accuracy to update one's credences by conditionalizing on the evidence.²¹ As noted in §2, the requirement to update by conditionalization is strictly stronger than Epistemic Conservativism for credences.

So when the relevant doxastic states are credence functions, Epistemic Conservativism follows immediately from Acc-Max and Strict Propriety. What about when the doxastic states are full beliefs? Here much hinges on the relation between full belief and credence. One attractive account of this relation comes from Lockeanism, the view that rational belief is a matter of having a credence above a certain threshold. Given Lockeanism, revising a full belief in p would require changing your credence in p, and we've just seen that changing your credence in p without new evidence would run afoul of Acc-Max, given Strict Propriety.

Of course, Lockeanism is controversial. However, Lockeanism can itself be justified from the point of view derived from epistemic decision theory. In particular, Easwaran 2016 and Dorst 2019 show that an agent maximizes expected accuracy just in case they believe all and only those propositions in which their credence exceeds a certain thresh-

²⁰See Oddie 1997; Greaves and Wallace 2006; Gibbard 2007; Joyce 2009; Horowitz 2013, for arguments in this vein. Some have worried that this argument only supports the weaker requirement that our scoring rule is *weakly proper*—that is, the expected accuracy of *c*, when calculated using *c*, is at least as great as the expected accuracy of any alternative credence function (e.g., Pettigrew 2011; Mayo-Wilson and Wheeler 2016). However, Campbell-Moore and Levinstein 2021, building on Shervish 1989, show that the requirement of weak propriety entails Strict Propriety given two further plausible constraints on a scoring rule. The first is *truth-directedness*: it says that the higher your credence in truths, and the lower your credence in falsehoods, the more accurate your credences are. The second is *additivity*: the accuracy of your credence function is defined over. The upshot: any truth-directed and additive scoring rule that is not self-undermining (and hence weakly proper) also satisfies Strict Propriety.

²¹Greaves and Wallace's argument assumes that a learning event can be represented as a partition of possibility space into a set of propositions that the agent might learn. See Schoenfield 2017 for arguments that this assumption will only hold in certain cases, and hence that the update procedure that maximizes expected accuracy is not updating by conditionalization on p, but rather on the proposition that one learned p. For other accuracy-based arguments for conditionalization, see Leitgeb and Pettigrew 2010; Easwaran 2013; Briggs and Pettigrew 2020.

old.²² [Can we dispense with Lockeanism altogether by defining an analogous notion of strict propriety that applies directly to full beliefs?]

4.3 Epistemic Decision Theory and Expert Deference

Epistemic decision theory is also in a position to underwrite expert deference principles. To see this, first we need to get clearer on the relevant notion of an 'expert.' Go back to our earlier example: you have no opinion about the weather tomorrow, but you know that your friend Ella, the esteemed meteorologist, does. From your point of view, Ella is an expert on whether it will rain tomorrow. But why? Here's a natural thought: because you expect Ella's opinions on this matter to be more likely to be accurate than yours.

This natural thought seems right, as far as it goes. But there is a wrinkle. As we have seen, there are different ways of scoring a credence function for accuracy. For this reason, Levinstein forthcoming proposes:

EXPERTISE AS ABSOLUTE SUPERIORITY A takes B to an expert with respect to some proposition p iff A expects B to be an absolute superior with respect to p—that is, A expects B's credence in p will be at least as accurate as B's according to every admissible scoring rule.

For Levinstein forthcoming, the admissible scoring rules are those that satisfy Strict Propriety, Truth-Directness, and Additivity.

Levinstein then shows that this definition entails that unless you expect the expert to have the very same credence in p as you do, you will expect the expert to be *strictly* more accurate according to every admissible scoring rule. That is:

EXPERTISE AS STRICT ABSOLUTE SUPERIORITY A takes B to an expert with respect to some proposition p iff A expects B to be a strict absolute superior with respect to p—that is, either A expects B's credence in p to be strictly more accurate than A's according to every admissible scoring rule, or A is sure that they both have the same credence in p.

Given this definition, Levinstein shows that it maximizes expected accuracy to satisfy a particular version of an expert deference principle for credences:

²²To get a sense of how it works, let T be the accuracy score of believing a truth, and let F be the accuracy score of believing a falsehood. One can then think of a believing p as a bet that pays off a positive score T in accuracy if you get p, and costs F if you get p wrong. From the point of view of expected accuracy, you should believe p if the expected benefit of believing truly is worth the expected cost of believing falsely. Easwaran and Dorst shows that if we assume (i) truth-directness for full beliefs (that is, the accuracy score of a belief depends on only on its truth-value), (ii) additivity for full beliefs (that is, the accuracy of an agent's full belief set is the the sum of the accuracy of their individual beliefs), and (iii) T > 0 > F, then it maximizes expected accuracy to believe p iff $c(p) \ge -F/(T - F)$.

SIMPLE TRUST Upon learning that an expert assigns a credence of at least n to p, a rational agent should assign a credence of at least n to p.²³

For the details, see Levinstein forthcoming. But to get an informal sense of how it works, suppose that right now you have a .5 credence in p: it will rain tomorrow. And suppose you think Ella is an expert with respect to p. Now, either you are certain that Ella's credence in p is the same as yours or you are not certain of this. Suppose you are certain of this. Then you are certain that Ella assigns .5 credence to p, so your current credence of .5 in p trivially satisfies Simple Trust. Suppose then you are not certain that Ella's credence in p is the same as yours. Then by Expertise as Strict Absolute Superiority, it follows that you expect Ella's credence in p to be strictly more accurate than your own. Now, if you do not know what Ella's credence in p is, there is nothing you can do with this expectation. Indeed, by Strict Propriety, you should still expect your current credence in p to be more accurate than any alternative credence function, and so you shouldn't revise it. However, suppose you learn that Ella assigns credence \geq .8 to p. Since you expect Ella's credence in p to be more accurate than your own, Acc-Max says you should revise your credence accordingly. Specifically, you should revise your credences so that you too assign \geq .8 credence to p, just as Simple Trust predicts.

So, given a natural way of defining the relevant notion of expertise, Simple Trust can be shown to maximize expected accuracy. To get from Simple Trust to Expert Deference, we can again appeal to Lockeanism, which, as we've seen, can itself be motivated by epistemic decision theory (§4.2). The conjunction of Simple Trust and Lockeanism entails Expert Deference.

Where does this leave us? In §2, I observed some striking parallels between the two forms of stare decisis and two plausible principles in epistemology: Epistemic Conservativism and Expert Deference. §3 went a step further, showing that, under certain conditions, we can derive the two forms of stare decisis from the corresponding epistemic principles. But this left a residual question: can we give a purely epistemic argument in support of Epistemic Conservativism and Expert Deference? This section answered in the affirmative; I showed that abiding by these norms will maximize expected accuracy. Combined with the argument in §3, this gives us an accuracy-based argument for *stare decisis*. Abiding by precedent is epistemically justified, because doing so will—at least under certain conditions—maximize the expected accuracy of a court's opinions.

This completes my epistemic argument for *stare decisis*. In what follows, I address some questions and objections.

²³Formally: let c be the credence function of a rational agent, and let c_E be a credence function that is an expert from the point of view of c. Then Simple Trust says that $c(p \mid c_E(p) \ge n) \ge n$. For further discussion and defense, see Dorst 2020; Dorst et al. 2021.

5 Q&A

Objection: 'In the law, it's common to distinguish between questions of fact and questions of law. When it comes to questions of fact, it makes perfect sense to talk about aiming at accuracy. But when it comes to questions of law, the whole idea of aiming at accuracy is misplaced.'

Reply: We should push back on the suggestion that we cannot meaningfully aim at accuracy on questions of law. Certainly some questions of law have objective answers. It is a question of law whether the U.S. Constitution allows for a two year old child to be the U.S. President. There is a clearly correct answer (ArtII.S1.C5.1). Of course, most questions of law that arise before courts are less straightforward. But this is not a reason to deny that we can sensibly apply the notions of accuracy and inaccuracy to these questions, any more than the difficulty of some scientific problem is a reason to deny that we can sensibly talk about which hypothesis provides an accurate explanation of the data.

What determines whether a given opinion on some matter of law is accurate or inaccurate? The answer will depend on the correct jurisprudence. Suppose a question of law hinges on the interpretation of some statute, and someone forms an opinion p on the matter. Textualists will hold that p is accurate just in case it accords with the best interpretation of the ordinary meanings of the words used in the statute. Intentionalists will hold that p is accurate just in case it corresponds with the legislature's intent in enacting the statute. Purposivists will claim the accuracy of p depends on the legislature's purpose in enacting the law. The present framework does not require settling between these different views. Whatever one's preferred jurisprudence, one can use it to derive accuracy conditions for opinions on questions of law.²⁴

Question: 'I'm sold on the idea that court opinions express beliefs. But why saddle courts with *credences*? This seems like an unhelpful idealization, if not a metaphysical extravagance.'

Reply: Judges certainly have credences, at least if credences are just degrees of belief. When it comes to many legal questions, their credences over the answer space should

²⁴While I am committed to claiming that we can meaningfully apply the notion of *accuracy* to questions of law, my framework leaves open the possibility that there are important differences between the metaphysical status of correct answers on questions of law and the metaphysical status of correct answers on questions of fact. It may be that the latter carve nature at its joints in a way that the former do not. For example, the framework developed here is perfectly compatible with expressivism about legal discourse, provided that our expressivism takes a 'quasi-realist' turn, allowing we can sensibly apply notions of truth and falsity to legal claims (cf. Blackburn 1993). A toy example: following Gibbard 2003 and Yalcin 2012, we could relativize accuracy to 'fact-prac' worlds, where a fact-prac world is an ordered pair of a world and some entity that represents a conative state of mind—say, a system of norms (Gibbard 1990), or a hyperplan (Gibbard 2003), or a normative perspective (Ridge 2014). We could then hold that when it comes to questions of fact, accuracy depends entirely on the fact parameter; when it comes to questions of law, accuracy also depends on the 'prac' parameter. All of our expected accuracy arguments would still run much as before.

not be all 1s and 0s. After all, the law is clearer on some matters than others. It is clear that the U.S. Constitution prohibits a two year old president. It is less clear whether the Second Amendment provides a right for citizens to carry firearms for reasons unconnected to militia service, or whether the administrative state is consistent with the non-delegation doctrine. A rational judge should recognize this, and apportion their degrees of confidence accordingly. This normative claim has just as much force when we are dealing not just with a single judge, but with a court.²⁵

When it comes to multi-member courts, there is an additional source of uncertainty. Even if every member of a court is certain of the answer to some legal question, they may disagree about what that answer is. In such cases, the court as a whole will often be less certain of the answer than its individual members. Consider the US Supreme Court's decision in *Furman v. Georgia*, where a highly fractured court ruled 5-4 that extant applications of the death penalty were unconstitutional, but could not agree why. It seems natural to describe the court as believing that the death penalty, as administered then, was unconstitutional. But it also seems natural to describe the court as less confident of this proposition than their unanimous opinions of the time.

The upshot: courts do—and should—have a range of levels of confidence in various legal questions. To model this, we need to ascribe more than just coarse-grained attitudes (belief, suspension, disbelief) to courts. Credences offer just what we need to model this fine-grained doxastic structure.

Question: 'You've told us that abiding by *stare decisis* maximizes expected accuracy. But sometimes a court should break with precedent. When should they do so, according to your framework?'

Reply: As formulated, Epistemic Conservatism just says that a rational agent shouldn't revise their doxastic states absent some positive reason to do so. It doesn't say what qualifies as a positive reason to do so, or how strong the reason needs to be. That said, epistemic decision theory suggests an answer. As noted in §4.2, Greaves and Wallace 2006 prove that on any strictly proper measure of accuracy, it always maximizes expected accuracy to revise your credences by conditionalizing on your evidence. Combined with Lockeanism, this result entails that an agent who maximizes expected accuracy will go from believing p to ceasing to believe p when—and only when—they acquire some new evidence e such that conditionalizing on e lowers their credence in p below the Lockean threshold. Similarly, they will go from believing p to believing $\neg p$ when—and only when—they acquire some new evidence e such that conditionalizing on e such that conditionalizing on e raises their

²⁵The idea that judges should sometimes be less than certain about some matters has not attracted much attention in legal scholarship, but the tides are slowly changing. Both Posner and Vermeule 2016 and Baude and Doerfler 2018 discuss the idea that judges should adjust their credences in light of peer disagreement (an idea I discuss further below). And in an important recent paper, Cox 2023 argues that judges should be less than fully certain of the correct jurisprudence, and this uncertainty will sometimes translate into rational uncertainty about the correct resolution of certain cases. This furnishes another argument for thinking that judges—and, by extension, courts—need to have credences defined over legal questions.

credence in $\neg p$ above the Lockean threshold. All this carries over straightforwardly to the courts.

Two points of elaboration. First, this gives us conditions under which a court epistemically ought to believe the negation of some precedent it formerly believed. From this it doesn't necessarily follow that in such circumstances a court ought to overrule its precedent. While believing $\neg p$ is plausibly a necessary condition on appropriately asserting $\neg p$, it is not sufficient. Presumably, it's not always appropriate to assert everything one believes! So further conditions may need to be be satisfied in order to make it appropriate to actively overrule a precedent. Some of these conditions may be epistemic—for example, perhaps appropriate assertion requires not just belief, but knowledge.²⁶ Others will be pragmatic; they will concern the expected consequences of making the assertion. Given the dependence on these pragmatic factors, we shouldn't expect to derive necessary and sufficient conditions for when a court should overrule a precedent from epistemic decision theory alone.

Second, this formal framework is silent on the substantive question of what sort of evidence might induce the relevant credal shifts on different legal questions. This just reflects a general feature of Bayesianism: what sorts of evidence will tip the balance from belief to disbelief will vary from agent to agent, depending on their priors. For example, a court that assigns a high prior credence to intentionalist approaches to legal interpretation will typically assign more weight to evidence of legislative intent (e.g., conference and committee reports, sponsor statements) than will a court with staunch textualist commitments.

Objection: 'Your argument for vertical *stare decisis* requires that lower courts regard higher courts as experts. Given Expertise as Strict Absolute Superiority (§4.3), this means that lower courts must either regard the higher courts as strictly more accurate than themselves on the matter at hand, or they must be sure that they and the higher court share the same credence on the matter at hand. But sometimes this condition is not satisfied. Sometimes a lower court may spot a clear flaw in a higher court's reasoning, or they may have more experience in cases on the topic at hand, or they may have reason to think that a higher court is corrupt, biased, ideologically motivated, etc.'

Reply: This is all true. And it reveals the limits of what we should expect from a purely epistemic argument for *stare decisis*. When a lower court is rationally convinced that a higher court is wrong on some point of law, or just generally incompetent, we shouldn't expect an accuracy-based argument for deferring to the higher court. In such cases, abiding by vertical *stare decisis* should be expected to result in *less* accurate decisions.

On reflection, this seems like the correct result. In such circumstances, we should expect the lower court to feel torn. On the one hand, purely epistemic considerations counsel against deference. On the other hand, there may still be strong pragmatic arguments in favor of deference—the very sort of pragmatic arguments that are widely

²⁶See fn.16.

discussed in the literature (§1). The lower court may well feel ambivalent about how to weigh these competing considerations.²⁷

There will also be situations where, from a purely epistemic point of view, a higher court ought to defer to a lower court. For example, if a lower court has considerably more experience in a specialized domain (say, the U.S. Tax Court), it may well be that a higher court should regard the lower court as strictly more accurate, and adjust its beliefs accordingly.

Objection: 'The sort of epistemic considerations you describe have little bearing on the actual practice of courts. When courts decide whether to stick with a precedent or give it up, they usually focus on pragmatic considerations.'

Reply: It's certainly right that the epistemic reasons for deferring to—or breaking with—precedent will not always align with the traditional pragmatic reasons. Indeed, we've just looked at one situation where these will come apart.

That said, we should reject the suggestion that epistemic considerations are irrelevant to the actual practice of courts, for two reasons. First, if evidence bearing on some precedent comes to light, court often will (and should) take that evidence into account when deciding whether to uphold the precedent. One of the most celebrated examples of overturning precedent comes from *Brown v. the Board of Education*, which held that the policy of "separate but equal" educational facilities violates the Equal Protection Clause of the 14th amendment, thereby overruling *Plessy v. Ferguson*. In their decision, the Warren court relied on evidence of the stigmatizing effect of segregated education, and the detrimental impact of such stigmatization on educational development. This evidence cast serious doubt on the empirical premises taken for granted in *Plessy*.

Second, evidential and pragmatic considerations are not mutually exclusive. In some circumstances, pragmatic considerations are themselves evidence for (or against) the accuracy of some judicial opinion.

Let me elaborate on this second point. In *Planned Parenthood of Southeastern Pa. v. Casey*, the U.S. Supreme Court laid out a number of considerations bearing on when a precedent should be overturned. These "Casey factors" include whether the precedent has proven unworkable, and whether factual circumstances have changed to render the precedent an "ill-fitting anachronism." Given the right set of priors, these pragmatic considerations could themselves be evidence bearing on the truth or falsity of the precedent. For example, suppose a court assigns a high prior to "living constitutionalism", the view that the correct interpretation of the constitution can evolve over time in light of changing social circumstances and moral norms. From this perspective, evidence that a previous interpretation of the constitution has proven unworkable or anachronistic could itself be evidence for the falsity of that interpretation. A similar point holds even for courts with originalist priors. Given the assumption that both the legislative intention and purpose of a statute is to provide workable guidance, then the fact that a

²⁷This sort of tradeoff between epistemic and practical concerns that has been amply discussed in other philosophical contexts (Pascal's wager, James' "will to believe"), but rarely in the legal domain.

previous interpretation has proved unworkable is evidence that previous interpretation of a statute does not fulfill either the legislative intention or the statute's purpose. And this will, in turn, be evidence that the previous interpretation of the constitution is *false*, according to intentionalist and purposivist views.

Objection: 'Your epistemic analysis of *stare decisis* overlooks a crucial feature of the role of precedent. Precedent *makes* law. The fact that a court decided some legal proposition p is part of what constitutes the truth of p.'

Reply: The idea that judges can "make" law through judicial fiat—sometimes known as the doctrine of "judicial discretion"—is very controversial. By now, gallons of ink have been spilled over whether judicial discretion is a genuine phenomenon, and, if so, whether it is a pervasive feature of the law.²⁸ In this paper, I have shown that even if one rejects the doctrine of judicial discretion, we can still give an epistemic vindication of *stare decisis*.

What if one accepts the doctrine? Then everything I've said in this paper still applies. It's just that we will have a much more direct argument that abiding by *stare decisis* maximizes the expected accuracy of a court's opinions. If precedent makes law, then judicial decisions are, at least under certain conditions, a form of self-verifying speech act: the act of announcing the decision makes its content true. An analogy: I am wondering whether Tim will speak today, so I ask him. Tim says, "Yes, I will." The fact that Tim made this assertion makes its content true. Since the fact that he made this assertion is now part of my evidence, it maximizes expected accuracy to believe its content.²⁹

Question: 'Suppose courts A and B are at the same level. A has set a precedent p some point of law, B has not ruled on this issue, but is inclined to rule $\neg p$. Should B be constrained by A's precedent?'

Reply: From the epistemic standpoint, this is most closely analogous to the problem of *peer disagreement* (Elga 2007). To illustrate with an example from Christensen 2007, suppose you and I are splitting the check. I do a quick calculation, and come up with the result that we each owe \$43. You announce that you got the result that we each owe \$45. How should your announcement affect my credence in how much we owe?

This question has sparked a lively debate. According to conciliationists (Elga 2007, Christensen 2007), I should give your answer equal weight to my own, and substantially reduce my credence that we each owe \$43. According to steadfasters (Kelly 2010), whichever party's credence was supported by the evidence is entitled to stick to their guns. For our purposes, we need not settle this debate here. The important point is that whatever we say about this issue should carry over to the analogous issue in the courts.

²⁸See e.g., Hart 1994; Dworkin 1963, 1967; Raz 1972; Greenawalt 1975; Hoffmaster 1982, among others.

²⁹A complication: if courts can also *remake* law through judicial fiat, then the act of overturning a precedent will change the legal facts. However, advocates of the doctrine of judicial discretion will presumably hold that there are some constraints on when judges can make (or remake) the law; not just anything goes. *Stare decisis* may itself be one of these constraints.

That said, epistemic decision theory offers to make progress on this issue in two respects. First, we've seen that epistemic decision theory can be used to show that updating by conditionalization maximizes expected accuracy. So this gives us, at a very general level, a recipe for how to respond to peer disagreement: update by conditionalizing on the proposition that you learn when you discover that you peer disagrees with you. Second, as Levinstein 2015 notes, a natural way of regimenting the notion of epistemic peerhood is to say that A regards B as an epistemic peer of C with respect to some topic iff A expects B and C's credences on that topic to be equally accurate. Given this definition, Levinstein shows that if you regard A and B as epistemic peers on some topic you should expect to give their credences on this topic equal weight. As a special case of this, if A regards B as an epistemic peer to A on some topic, then A should expect to assign B's credence on that topic as much weight as A's own credence. This highlights another way in which the tools of epistemology can be brought to bear on questions of when to defer to precedent.³⁰

Question: 'Suppose at t_0 , p is precedent. But at t_1 , a court overrules this precedent, holding $\neg p$. At t_2 , the same court is faced with a new case that raises the same point of law. Should they defer to the original precedent (p) or the new one $(\neg p)$?'

Reply: An initial comment: put this way, the question might sound *recherché*, but it has real world urgency. The majority opinion in *Dobbs* in effect claimed that the history of abortion rights in the U.S. takes this shape. One of Alito's central arguments for overruling *Roe* was that *Roe* had broken with precedent, and was therefore not itself entitled to the deference usually accorded to precedent. So the reproductive rights of millions of Americans hinged on a question of precisely this form.

How should we analyze this sort of case from the epistemic point of view? The answer is rather nuanced. But for starters we should note that at t_1 , the court—if sincere believes $\neg p$. So if they receive no new relevant evidence between t_1 and t_2 , they should retain this belief (by Epistemic Conservativism). From their perspective at t_2 , abandoning their belief in $\neg p$ should be expected to lead to a loss in accuracy (given Strict Propriety and Lockeanism).

³⁰For the most part, the topic of peer disagreement has not received much attention in legal scholarship, though see Posner and Vermeule 2016, and, esp., Baude and Doerfler 2018, both of which advocate some form of conciliationism in the face of judicial disagreement. However, neither of these papers embeds this question in the framework of epistemic decision theory, or the more general goal of maximizing expected accuracy. This has consequences for some of their specific claims. For example, Baude and Doerfler 2018 suggest that "two judges ought to consider one another 'epistemic peers' only to the extent that they share the same judicial outlook or methodology" (p.326). But from the perspective of epistemic decision theory, shared judicial outlook is neither necessary nor sufficient for epistemic peerhood. It is certainly not sufficient, since A and B might share a judicial outlook but one of them has evidence that the other lacks. And it is not necessary, because I might regard A and B as equally accurate on some legal question based on their track record, even if I know nothing about their judicial outlook, or even if I know for a fact that they have divergent outlooks. In such a case, I should still expect to give their credences on this question equal weight.

However, there is a wrinkle. We sometimes revise our beliefs without being aware that we are doing so. At t_1 , the court might not have been aware that they were breaking with precedent. Suppose that at t_2 , the court discovers this fact. That is, they now learn e: When we held $\neg p$ at t_1 , we were breaking with precedent. This fact is now new evidence that the court possesses at t_2 . So maximizing expected accuracy requires conditionalizing on this evidence, which may affect the court's credence in $\neg p$. Whether it does so will all depend on (i) the court's prior conditional credences, (ii) whether between t_1 and t_2 the court gained any further evidence e' in addition to e—for example, evidence about the reasons for which the court broke with precedent at t_1 , or any other evidence that bears on the truth of p. If so, then maximizing expected accuracy requires taking this further evidence into account; even if $c(\neg p|e) < c(\neg p)$, it might be that $c(\neg p|e \land e') \ge c(\neg p)$. Given the dependence on such factors, there is no general, one-size-fits-all answer to our question. That said, viewing this issue through the lens of epistemic decision theory is still fruitful, because it helps think systematically about which features impact the accuracy of the court's opinion and how they do so.

References

William Baude and Ryan Doerfler. Arguing with friends. *Michigan law Review*, 117, 2018. Simon Blackburn. *Essays in Quasi-Realism*. Oxford University Press, Oxford, 1993.

- Rene Bollinger. Explaining the justificatory asymmetry between statistical and individualized evidence. In *The Social Epistemology of Legal Trials*, pages 60–76. Routledge, forthcoming.
- R.A. Briggs and Richard Pettigrew. An accuracy-dominance argument for conditionalization. *Noûs*, 54(1):162–181, 2020.
- Lara Buchak. Belief, credence, and norms. Philosophical Studies, 169:285-311, 2014.
- Steven Burton. The conflict between stare decisis and overruling in constitutional adjudication. *Cardozo Law Review*, 35(5), 2014.
- Catrin Campbell-Moore and Benjamin A Levinstein. Strict propriety is weak. *Analysis*, 81(1): 8–13, 2021.
- Roderick Chisholm. A version of foundationalism. *Midwest Studies in Philosophy*, 5(1):543–564, 1980.
- David Christensen. Conservatism in epistemology. Noûs, 28(1):69-89, 1994.
- David Christensen. Putting Logic in its Place. Oxford University Press, 2005.
- David Christensen. Epistemology of disagreement: the good news. *Philosophical Review*, 116(2): 187–217, 2007.
- Courtney Cox. The uncertain judge. University of Chicago Law Review, 90(3), 2023.
- Kevin Dorst. Lockeans maximize expected accuracy. Mind, 128(509):175-211, 2019.
- Kevin Dorst. Evidence: A guide for the uncertain. *Philosophy and Phenomenological Research*, 100(3):586–632, 2020.
- Kevin Dorst, Benjamin A Levinstein, Bernhard Salow, Brooke E. Husac, and Branden Fitelson. Deference done better. *Philosophical Perspectives*, 35(1):99–150, 2021.
- Neil Duxbury. *The Nature and Authority of Precedent*. Cambridge University Press, Cambridge, 2008.
- Ronald Dworkin. Judicial discretion. Journal of Philosophy, 60:624–638, 1963.
- Ronald Dworkin. The model of rules I. University of Chicago Law Review, 35(1), 1967.

- Kenny Easwaran. Expected accuracy supports conditionalization and conglomerability and reflection'. *Philosophy of Science*, 80(1):119–142, 2013.
- Kenny Easwaran. Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Noûs*, 50(4):816–853, 2016.
- Adam Elga. Reflection and disagreement. Noûs, 41(3):478-502, 2007.
- David Enoch, Levi Spectre, and Talia Fisher. Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy and Public Affairs*, 40(3):197–224, 2012.
- Richard Foley. Epistemic conservativism. Philosophical Studies, 43(2):165-182, 1983.
- Richard Foley. Working Without a Net. Oxford University Press, 1993.
- Peter A. French. *Collective and Corporate Responsibility*. Columbia University Press, New York, 1984.
- Allan Gibbard. Wise Choices, Apt Feelings. Harvard University Press, Cambridge, MA, 1990.
- Allan Gibbard. Thinking How to Live. Harvard University Press, Cambridge, MA, 2003.
- Allan Gibbard. Rational credence and the value of truth. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 2. Oxford University Press, New York, 2007.
- Arthur Goodhart. Precedent in english and continental law. L.Q. Review, 40(58), 1934.
- Hilary Greaves and David Wallace. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632, 2006.
- Kent Greenawalt. Discretion and judicial decision: The elusive quest for the fetters that bind judges. *Columbia Law Review*, 75, 1975.
- Kent Greenawalt. Policy, rights, and judicial decision. *Georgia Law Review*, 11(5):991–1054, 1977. H.L.A. Hart. *The Concept of Law*. Clarendon Press, Oxford, 2 edition, 1994.
- Barry Hoffmaster. Understanding judicial discretion. Law and Philosophy, 1(1):21–55, 1982.
- Sophie Horowitz. Immoderately rational. Philosophical Studies, 167(1):1-16, 2013.
- Michael Huemer. Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74:30–55, 2007.
- James Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.
- James Joyce. Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In *Degrees of Belief*, pages 263–297. Springer, 2009.
- Thomas Kelly. Peer disagreement and higher-order evidence. In Feldman and Warfield, editors, *Disagreement*. Oxford University PRess, Oxford, 2010.
- Lewis A. Kornhauser. An economic perspective on stare decisis. *Chicago-Kent Law Review*, 65 (1), 1989.
- Hannes Leitgeb. The Stability Theory of Belief. Philosophical Review, 123(3):131-171, 2014.
- Hannes Leitgeb and Richard Pettigrew. An objective justification of bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*, 77(2):236–272, 2010.
- Benjamin A Levinstein. With all due respect: The macro-epistemology of disagreement. *Philosophers' Imprint*, 15(13), 2015.

Benjamin A Levinstein. Accuracy, deference, and chance. The Philosophical Review, forthcoming.

David Lewis. A Subjectivist's Guide to Objective Chance. In Jeffrey, editor, *Studies in Inductive Logic and Probability*, volume 2, pages 263–293. University of California Press, 1980.

Christian List. What is it like to be a group agent? *Noûs*, 52(2):295–319, 2018.

Christian List and Philip Pettit. Group Agency: The Possibility, Design, and Status of Corporate Agents. Oxford University Press, Oxford, 2011.

W.G. Lycan. Judgement and Justification. Cambridge University Press, Cambridge, 1988.

- Conor Mayo-Wilson and Gregory Wheeler. Scoring imprecise credences: A mildly immodest proposal. *Philosophy and Phenomenological Research*, 93(1):55–78, 2016.
- Kevin McCain. The virtues of epistemic conservatism. Synthese, 164(2):185-200, 2008.
- G.E. Moore. Commonplace Book: 1919-1953. Allen and Unwin, London, 1962.
- Sarah Moss. Scoring rules and epistemic compromise. Mind, 120(480):1053-1069, 2011.
- Graham Oddie. Conditionalization, Cogency, and Cognitive Value. British Journal for the Philosophy of Science, 48:533-541, 1997.
- Richard Pettigrew. An improper introduction to epistemic utility theory. pages 287–301. Amsterdam, 2011.
- Richard Pettigrew. Accuracy and the Laws of Credence. Oxford University Press, Oxford, 2016.
- Eric Posner and Adrian Vermeule. The votes of other judges. *Georgetown Law Journal*, 105, 2016. James Pryor. The skeptic and the dogmatist. *Noûs*, 34(4):517–549, 2000.
- W.V.O. Quine. Two dogmas of empiricism. Philosophical Review, 60(1):20-43, 1980.
- John Rawls. A Theory of Justice. Harvard University Press, Cambridge, 1971.
- Joseph Raz. Legal principles and the limits of law. Yale Law Journal, 81, 1972.
- Michael Ridge. Impassioned Belief. Oxford University Press, New York, 2014.
- Lewis Ross. Rehabilitating statistical evidence. *Philosophy and Phenomenological Research*, 102 (1):3–23, 2019.
- Carol Anne Rovane. The Bounds of Agency: An Essay in Revisionary Metaphysics. Princeton University Press, Princeton, 1997.
- Frederick Schauer. Precedent. Stanford Law Review, 39(3):571-605, 1987.
- Miriam Schoenfield. Bridging rationality and accuracy. *Journal of Philosophy*, 112(12):633–657, 2016.
- Miriam Schoenfield. Conditionalization does not (in general) maximize expected accuracy. *Mind*, 126(504):1155–1187, 2017.
- Mark Shervish. A general method for comparing probability assessors. *The Annals of Statistics*, 17, 1989.
- Lawrence Sklar. Methodological conservatisim. Philosophical Review, 84(3):374-400, 1975.
- Martin Smith. When does evidence suffice for conviction? Mind, 127(508):1193-1218, 2018.
- John Paul Stevens. The life span of a judge-made rule. NYU Law Review, 58(1), 1983.
- Scott Sturgeon. Reason and the Grain of Belief. Noûs, 42(1):359–396, 2008.

Deborah Tollefsen. Groups as Agents. Polity Press, Cambridge, 2015.

Laurence Tribe. Trial by mathematics: Precision and ritual in the legal profession. *Harvard Law Review*, 84(6), 1971.

Peter Unger. Ignorance. Clarendon Press, Oxford, 1975.

Bas van Fraassen. Belief and the will. Journal of Philosophy, 81:235-256, 1984.

- Jeremy Waldron. Stare decisis and the rule of law: A layered approach. *Michigan Law Review*, 111(1), 2012.
- Ryan Wasserstrom. The Judicial Decision. Stanford University Press, Stanford, 1961.

Timothy Williamson. Knowledge and its Limits. Oxford University Press, Oxford, 2000.

Seth Yalcin. Bayesian Expressivism. Proceedings of the Aristotelian Society, 112(2):123–160, 2012.