

Gradational Accuracy and Non-classical semantics

J. Robert G. Williams

Last revised: 8th February 2011

Abstract

This paper gives a generalization of Jim Joyce’s 1998 argument for probabilism, dropping his background assumption that logic and semantics are classical. Given a wide variety of non-classical truth-value assignments, Joyce-style arguments go through, allowing us to identify in each case a class of “non-classically coherent” belief states. To give a local characterization of coherence, we need to identify a notion of logical consequence to use in an axiomization. There is a very general, ‘no drop in truth value’ characterization that will do the job. The result complements Paris’s 2001 discussion of generalized forms of dutch books appropriate to non-classical settings.

Contents

1	The accuracy argument	2
2	Revisionary semantics, probability and logic	3
3	Non-classical truth values	5
4	Gradational accuracy	9
5	The generalized proof	10
6	Domination immunity	12
7	Axiomatizing coherent credences	15

Joyce (1998) gives an argument for probabilism: the doctrine that rational credences should conform to the axioms of probability. In doing so, he provides a distinctive take on how the normative force of probabilism relates to the injunction to believe what is true. But Joyce presupposes that the truth values of the propositions over which credences are defined are classical. I generalize the core of Joyce’s argument to remove this presupposition. On the same assumptions as Joyce uses, the credences of a rational agent should always be *weighted averages of truth value assignments*. In the special case where the truth values are classical, the weighted averages of truth value assignments are exactly the probability functions. In the more general case, probabilistic axioms formulated in terms of classical logic are violated—but we will show that generalized versions of the axioms formulated in terms of non-classical logics are satisfied.

The plan of campaign is as follows. **Section 1** describes the pattern of Joyce’s original argument for probabilism, and locates the component of the argument under discussion in this paper. **Section 2** identifies the way in which Joyce’s argument presupposes a classical setting, and motivates interest in non-classical generalizations. In **section 3** I set out a smorgasbord of non-classical conceptions of truth-values, and characterize an accompanying logic for each one. **Sections 4 and 5** describe Joyce’s argument and its generalization. **Section 6** considers whether a converse to Joyce’s result is necessary for it to have philosophical significance. **Section 7** considers axiomatizations of the convex sets of truth values.

1 The accuracy argument

Probabilism is the doctrine that one’s degrees of belief should meet the constraints of the probability calculus. Admittedly, we fall short of this ideal. According to probabilism, such failures are violations of ideal coherence—flaws to be removed.

One important question is whether probabilism is true: sceptics

might wonder whether there are any general coherence norms on belief of this level of generality; and even if they think such norms exist, they might think they are looser than probabilists allow. But even if we accept probabilism itself pro tem, there’s more work to do; to explain *why* probabilistic credences are something to aspire to. Aiming to have true beliefs, or beliefs that count as knowledge, might be explanatory bedrock; but it seems odd to take as basic the aspiration to achieve a certain pattern among our cognitive states. It would be satisfying to be able to explain the normative punch of probabilism, in terms of aspects of beliefs we care about independently (cf. Kolodny, 2007; Broome, 2005).

Joyce (1998) argues that insofar as we care about degrees of belief being *accurate* (as close as possible to the truth), we can state precisely the way that lapses from the probabilistic ideal are a *flaw*.¹ On his assumptions, Joyce shows that if a belief state b violates the probability axioms, it will be *dominated* by an alternative belief state c that meets the probability axioms, in the sense that one can know, of c , that no matter what the world is like, c is more accurate (hence, epistemically ‘better’) than b . The pivotal assumption is that there’s some privileged way of measuring the “inaccuracy” of degrees of belief, that epistemically virtuous agents should minimize. Joyce argues that this satisfies certain constraints. Now suppose Sally has degrees of belief c that are not structured probabilistically. Joyce’s theorem shows that Sally is in a position to know, a priori, of some specific belief state b , that it is more accurate than her own according to the One True measure of (in)accuracy. An agent aware of this fact should not stick with her original, dominated beliefs, but rather move to the more accurate ones. Even if Sally doesn’t actually become aware of this fact, it remains the case that her beliefs are underminable by pure reflection—and this kind of instability is the kind of flaw that

¹I am taking the argument in this early paper as my model. NN and MM have suggested to me that in Joyce (2009) differs not only in the formal proofs it offers, but also in its conception of the normative role of accuracy.

should not be present in an ideally rational believer. This, I take it, constitutes the heart of the Joycean explanation of why probabilism is an appropriate ideal.

Even if this explanation of the normative punch of probabilism is successful, it is a separate question whether it should convince the previously unconvinced of the truth of probabilism. This will turn on whether there is *independent* warrant for taking the Joycean axioms as true of the accuracy measure.² If we're concerned with the former project, however, potentially probabilism itself could feature as a premise in a justification of Joyce's axioms.

There are three steps to Joyce's argument, once the general idea of an accuracy norm on degrees of belief is accepted. The first step is an articulation of axiomatic constraints on accuracy, and justifications thereof. The second is to prove the accuracy-domination theorem, relative to those constraints. The third is to use that theorem, in something like the way sketched above, to explain or argue for probabilism. In what follows, I focus on the second step, *presupposing* the Joycean axiomatic constraints, and *assuming* that results if successful have a bearing on probabilism (though one particular worry will be discussed in detail later).

2 Revisionary semantics, probability and logic

Joyce's original setting presupposed that the objects of credence behave classically. Thinking of them as sentences in a 'language of thought', the assumption is that relative to each possible situation, the sentence is either true, or false and the overall distribution meets

²Joyce (1998) provides such arguments. The original justifications were a little thin in places; see (Maher, 2002). Joyce (2004) sketches a promising rationale for the two most contentious axioms, based on the assumption that certain 'uniform' probability distributions are epistemically permissible. There are alternative axiomatic bases to be considered; see Gibbard (2008) for scepticism about this project.

the familiar classical requirements (analogously, if credences attach to propositions, the assumption is that propositions map each possible world to one of the two classical truth values). But there are many who argue that this assumption fails, in general or for certain subject matters. For example, perhaps cases of presupposition failure ("The King of France is bald") are neither true nor false.³ Perhaps observational predicates obey an intuitionistic, rather than classical, calculus.⁴ Perhaps it's neither true nor false to say that Newtonian mass is relativistic mass.⁵ Perhaps the law of excluded middle should be given up in order to get a satisfying take on semantic paradoxes.⁶ Perhaps borderline cases of vague predicates require us to think about degrees of truth, or buy into a supervaluational framework.⁷

Semantically revisionary theories such as these are often given without an *explicit* description of the implications for rational belief.⁸ This is unfortunate. Importantly different revisionary theories are not distinguished; the characterisation of truth and logic is often treated as if it was a merely terminological issue rather than theoretically central; and the wider ramifications of a revisionary theory of belief—for example, in the theory of rational decision—are obscured. As Williamson (2006) urges, those who propose to alter the foundations on which much successful work has been faced (a case in point being classical logic and semantics as an underlying assumption of decision theory) are under an obligation to show how that work may be reconstructed. Explaining how constraints on rational belief adapt to the revised setting is a prerequisite.

This is where the accuracy-rationale for probabilism will assist. As we will see, the Joyce-style explanation of probabilistic norms is suffi-

³This thought goes back to Strawson.

⁴Wright (2003).

⁵Field (1973).

⁶See Field (2008), Maudlin (2004) for two recent (and very different) suggestions along these lines.

⁷See Machina (1976); Smith (2008) and Keefe (2000); Fine (1975), respectively.

⁸Exceptions include Field (2003), Priest (2006), Smith (2008).

ciently modular that we can remove the specifically classical elements, switch in non-classical replacements, and run the argument again. We get a *principled* rationale for a particular (nonclassical) ideal for rational belief, relative to each non-classical setting.⁹

A direct benefit from this approach is to allow the non-classicist a distinctive line on how to generalize from what she says about semantics (and truth-value distributions) to what she should say about logic. Typically, if one has more than the classicist’s two truth values to play with, there will be more than one way to generalize the classicist’s characterization of consequence. We could perhaps trust to putative “platitudes” such as “logical consequence is necessary truth preservation”, but this runs into a string of difficulties: it’s not clear whether the platitude should really be trusted in this context, nor why that which satisfies such platitudes should be more theoretically interesting than other related notions; and in any case, the ‘platitude’ only tells us what consequence is if we’ve already singled out one of the non-classical truth values as deserving the name “truth”. We can do much better. One thing that the present results deliver is a reason to be interested in *one of these characterizations in particular*. For by the accuracy-domination arguments, beliefs should be non-classical probabilities, and just as classical logic can be used to axiomatize classical probabilities, non-classical logic (if given a particular characterization) can be used to axiomatize non-classical probabilities. In short, one particular characterization of logic in a given revised setting generalizes the belief-norming role of classical consequence.

One might worry that our ambition for the generalized Joyce-style argument is overly strong, and makes the enterprise dubious. After all, don’t the paracomplete (Kleene-3) and the paraconsistent (LP)

⁹Our results here are directly comparable to those of Paris (2001), who uses a generalized form of Dutch Book argument to motivate particular generalizations of probability theory in the non-classical setting. The core result Paris proves is that any credence function that isn’t a convex mixture of truth-value distributions is Dutch-Bookable, in a particular sense. Here we show that not being a convex mixture is sufficient for accuracy-domination.

reading of one and the same strong Kleene truth tables correspond to *different* non-classical views? And if so, shouldn’t we expect a plurality of views on rational belief, logic and the rest, rather than being straightjacketed into a single choice? All we need to make room for this is that observation that the Kleene truth tables alone do not fix the interpretation of the ‘truth statuses’ (let’s call them, neutrally, *A*, *B* and *C*) that fill their cells—indeed, unless we can somehow independently distinguish a truth table from its dual, then we won’t even be in a position to pick out which of these is most ‘truth-like’ and which is most ‘falsity-like’.

I suggest that what is required is that we specify a value for each status that describes how it functions as an *aim for belief*. There are lots of prima facie coherent ways of assigning such values to the Kleene tables. Thus, one view would have it that, from a God’s eye perspective, one should have credence 1 in any sentence that has status *A*; and credence 0 in a sentence that has status *B* or *C*. A different view has it that, one should have credence 1 in sentences with status *A* or *B*, and otherwise credence 0. A third view agrees with the other two on *A* and *C*, but insists that one should have credence 0.5 in any sentence with credence *B*. These ideal aims for credence I will call truth *values*, and I assume that a non-classical (or, indeed, a classical) semantics is incomplete without such a specification.

Once we add distinctive value-assignments that a paraconsistentist or paracompleteist is likely to favour, one and the same general characterization of consequence (‘no drop in truth value’) returns a logic; but which logic it is will vary depending on what one said about the values of the different statuses (in one case it might be paracomplete, in another paraconsistent, etc). This is the logic which will turn out to have the central role in norming belief.¹⁰

¹⁰Even if one isn’t a convinced non-classicist, one should be interested in these results. Presumably the principled agnostic will have credences that are a mix between appropriate to a convinced classicist and those of the convinced non-classicist. This too can be given a rationale in this framework. In the framework

3 Non-classical truth values

The previous section argued that to fully describe a non-classical semantics, we need not only a specification of truth-statuses, and their compositional interrelations, but also how such statuses function as aims for credence. This section lays out a smorgasboard of non-classical semantics, taking care in each case to specify the truth *value*. Often theories with essentially the same kind of underlying structure differ markedly at the level of truth value assignments—and as we’ll see, this will have substantial knock-on consequences. Once we have our range of theories to play with, we’ll reprove the Joyce accuracy domination argument (with accuracy now a matter of distance from the nonclassical truth values); and look at what patterns among the rational belief states arise.

For starters, W is a (finite) space of worlds.¹¹ Ω is a set of ‘propositions’. Ω is equipped with a structure. For any $p, q \in \Omega$, there are propositions $p \wedge q$, $\neg p$ and $p \vee q \in \Omega$. Think of \wedge , \neg and \vee as functions mapping (/pairs of) propositions to propositions. Where appropriate, we’ll also include a \rightarrow with the same status.¹² (I will exploit

to be developed below, generalized probabilities are defended relative to a certain class of ‘possible truth value assignments’. The question is: what truth-value assignments count as possible? If classicism is universally true, all genuinely possible truth-value assignments will be classical. The friend of non-classicism has a more liberal view. The agnostic may be agnostic about which truth-value assignments are genuinely possible—but she needn’t give up on the framework. Non-classical truth-value assignments are *epistemically* possible for her, and we can run the arguments below relative to her *epistemic* possibilities. This adaption of the arguments will allow us to formulate constraints that *even by the lights of the possibilities she treats as epistemically open*, substantial (ideal) constraints on her credences are in force.

¹¹The restriction to finite world-spaces is potentially significant—one might reasonably think that the belief states of a real-world agent would require infinitely many possibilities. Note that this restriction isn’t appealed to in Joyce’s proof, though many of the example scoring rules satisfying his axioms only make sense over finite spaces.

¹²This is important in settings where conditionals are not definable in terms of

Russell-style equivocation on the nature of these entities. You can think of these as Fregean thoughts, or sentences of a natural language or a ‘language of thought’, etc). We interpret the propositions via the function $i : \Omega \times W \rightarrow S$, where S is a set of ‘statuses’. We write $p(w)$ for $i(p, w)$, for fixed interpretation i .¹³

We also define a map $\llbracket \bullet \rrbracket : S \rightarrow [0, 1]$, which gives us the *truth value* of the status S . $\llbracket p(w) \rrbracket$, which for ease of notation I will write as $\llbracket p \rrbracket_w$, then gives the *truth value* of the proposition p at world w .

This description leaves open numerous issues. Most pressingly, it leaves open what S is, and what $\llbracket \bullet \rrbracket$ should be. Connectedly, we haven’t said how \wedge , \neg and \rightarrow should interact with the assignment of statuses and truth values. How we fill in the picture will fix what kind of semantics (classical or non-classical) we are offering. To save space, I’ll give the clauses for \wedge and \neg , and unless otherwise stated $p \vee q$ will take the same status as $\neg(\neg p \wedge \neg q)$ and $p \rightarrow q$ will take the same status as $\neg p \vee q$.

Classical $S = \{\text{true}, \text{false}\}$; $\llbracket \bullet \rrbracket$ sends true to 1 and false to 0. $(p \wedge q)(w) \mapsto \text{true}$ iff $p(w) \mapsto \text{true}$ and $q(w) \mapsto \text{true}$. $\neg p(w) \mapsto \text{true}$ iff $p(w) \mapsto \text{false}$.

Kleene gaps $S = \{\text{true}, \text{neither}, \text{false}\}$; $\llbracket \bullet \rrbracket$ sends true to 1 and other statuses to 0. $(p \wedge q)(w) \mapsto \text{true}$ iff $p(w) \mapsto \text{true}$ and $q(w) \mapsto \text{true}$.

disjunction and negation. The Łukasiewicz logics are good examples of this, and the expressive power of tautologies involving the conditional are important to the later results.

¹³An alternative would be to think of the propositions as having their interpretations essentially—we could then just identify them with functions from worlds into S . The only difference this would make for us is that we’d have to rethink the characterization of logic, below, which as is traditional is given in terms of varying interpretations. One option is to expand the space W so it contains enough worlds to induce arbitrary truth value assignments, and define consequence by generalizing over worlds. However, we’ve assumed that W is finite, which means there just won’t be enough on them. So unless the results can be generalized to remove the assumption of finiteness, the methodology adopted here seems preferable.

$(p \wedge q)(w) \mapsto \text{false}$ iff either $p(w) \mapsto \text{false}$ or $q(w) \mapsto \text{false}$. Otherwise $(p \wedge q)(w) \mapsto \text{neither}$. $\neg p(w) \mapsto \text{true, neither, false}$ iff $p(w) \mapsto \text{false, neither, true}$ respectively.¹⁴

LP gluts $S = \{\text{true, both, false}\}$; $\llbracket \bullet \rrbracket$ sends **true, both** to 1 and **false** to 0. $(p \wedge q)(w) \mapsto \text{true}$ iff $p(w) \mapsto \text{true}$ and $q(w) \mapsto \text{true}$. $(p \wedge q)(w) \mapsto \text{false}$ iff either $p(w) \mapsto \text{false}$ or $q(w) \mapsto \text{false}$. Otherwise $(p \wedge q)(w) \mapsto \text{both}$. $\neg p(w) \mapsto \text{true, both, false}$ iff $p(w) \mapsto \text{false, both, true}$ respectively. (Aside from the labelling, the only difference between this and Kleene gaps is the projection from statuses to truth values).¹⁵

Fuzzy $S = [1, 0]$ and $\llbracket \bullet \rrbracket$ is identity. $(p \wedge q)(w) \mapsto \min(p(w), q(w))$. $(p \vee q)(w) \mapsto \max(p(w), q(w))$. $\neg p(w) = 1 - p(w)$. $(p \rightarrow q)(w) \mapsto 1 - (p(w) - q(w))$, if $p(w) \geq q(w)$, and otherwise is set to 1.¹⁶

Finite fuzzy Exactly as above, except that for some choice of n , $S = \{m/n : 0 \leq m \leq n\}$.

Fuzzy gaps (finite or infinite) As for the above two cases, except $\llbracket \bullet \rrbracket$ maps 1 to 1, and any element of $(1, 0]$ to 0.¹⁷

Supervaluations S is a set of functions from ‘delineations’ to $\{\text{true, false}\}$. In connection to vagueness, David Lewis (1970) thinks of a delineation as a ‘sequence of boundary-specifying numbers’—one coordinate of which, for example, would mark a possible boundary between things that count as warm and things that count as cool—the truth-value of the proposition that the plate is warm at a delineation would turn on whether it’s temperature was above or below the boundary point the delineation specifies. Other conceptions are possible: on a Thomason-style model of

future contingents sentences, the delineations may be identified with maximal linear big-bang-to-heat-death world histories, with truth values over claims being induced in the natural way.

Whatever we say about the nature of the delineations, we require the overall pattern of truth values induced by a delineation behave classically in the following sense: $(p \wedge q)(w)(x) = \text{true}$ iff $p(w)(x)$ and $q(w)(x) = \text{true}$. $\neg p(w)(x) = \text{true}$ iff $p(w)(x) = \text{false}$. So each delineation induces a classical truth-value assignment, in the sense given earlier.

Finally, there is the question of extracting truth values. To this end, we need to appeal to a privileged subset of these delineations—the *admissible* ones. For the vagueness case, these may be those delineations which induce classical truth-value assignments over sentences that are *consistent with the meaning-fixing facts*. For the case of future contingents, relative to a given moment in time, the admissible delineations may be those world-histories which are ‘historically possible’ at that moment, e.g. consistent with the laws of nature and the actual history up to that point in time. However they are picked out, we say that $\llbracket s \rrbracket = 1$ iff s is the function which maps every admissible delineation to **true**. Otherwise $\llbracket s \rrbracket = 0$. (Having value 1, by this definition, corresponds to being *supertrue* in the standard supervaluationist lingo).¹⁸

Degree supervaluations Exactly as above, except when it comes to extracting truth values. Here, $\llbracket s \rrbracket = d$ iff s maps a proportion d of the delineations to **true**, and the rest to the **false** (more generally, relative to a measure over the space of delineations, $\llbracket s \rrbracket$ will be the measure of the subset of delineations that s maps to **true**).¹⁹

Intuitionism Let S be a set of pairs of Kripke structures²⁰ and as-

¹⁴Cf. e.g. Priest (2001, ch.7)

¹⁵Cf. e.g. Priest (2001, ch.7)

¹⁶Note that the classical setting above is a (notational variant of the) limiting case of this fuzzy setting, i.e. the case where all sentences get values from $\{1, 0\}$ rather than drawn from the whole of $[1, 0]$. Cf. e.g. Priest (2001, ch.11)

¹⁷Cf. e.g. Priest (2001, ch.7,11)

¹⁸For more in this semantics, cf. e.g. Varzi (2007).

¹⁹For more on this semantics, cf. Lewis (1970); Kamp (1975); Edgington (1997).

²⁰That is, a partially ordered Kripke frame $\langle K, \leq \rangle$, where K is a non-empty set

signments of exactly one of truth, falsity to each node of that structure.

We shall assume that each world w ‘comes equipped’ with a Kripke structure $K(w)$, such that for any proposition p , $p(w)$ always takes a value whose first element is $K(w)$.

We write $p(w)(n)$ for the value assigned to node n of $K(w)$ by the relevant assignment function.²¹

We take as basic a distinction between atomic and non-atomic propositions.²² For atomic propositions, we insist that the assignment of truth values to nodes is *persistent*.²³ We then require that $(p \wedge q)(w)(n)$ is true iff both $p(w)(n)$ and $q(w)(n)$ are true; $(p \vee q)(w)(n)$ is true iff either $p(w)(n)$ or $q(w)(n)$ are true; that $\neg p(w)(n)$ is true iff for every m below n in $K(w)$, there’s a o below m where $p(w)(o)$ is false; $(p \rightarrow q)(w)(n)$ is true iff for every m below n in $K(w)$ such that $p(w)(m)$ is true, we have $q(w)(m)$ true. We then say that $p(w)$ is forced if σ_w maps each node of $K(w)$ to true.

To define the truth value assignment, we replicate the definition of forcing. We say that $\llbracket p \rrbracket_w = 1$ iff $p(w)(n)$ is true for every node n of $K(w)$. Otherwise $\llbracket p \rrbracket_w = 0$.²⁴

As anticipated earlier, the main role in this paper for the projection

(the elements of which we call ‘nodes’) and \leq is a partial ordering of the set.

²¹If we define a relation \vDash_w between nodes and propositions (relative to a choice of world) such that $n \vDash_w p$ iff $p(w)(n) = \text{true}$, then the intention is that $\langle K(w), \vDash_w \rangle$ form a Kripke model for intuitionistic logic in the usual sense.

²²This is the one time in this preamble that we need to assume that we can identify atomic propositions, from which all other propositions are built via the standard connectives.

²³An assignment σ is persistent if when σ assigns true to n , then it assigns the same to all nodes below n .

²⁴For more on intuitionism, cf. e.g. Priest (2001, ch 6)

from ‘statuses’ to ‘values’ in $[1,0]$ will be as *aims for belief*.²⁵ Generalizing Joyce’s idea of ‘accuracy’ from the classical case, we will be assuming that credences are *better* the closer they are to the truth values (on the intended interpretation). Suppose, for example, that one thought that semantic indeterminacy is appropriately represented via supervaluational semantics, and that ‘Harry is bald’ is borderline—true on some admissible ‘delineations’ of the terms, and false on others. This idea is represented by a semantics assigning statuses to propositions of the form common to ‘supervaluational’ and ‘degree supervaluational’ systems above. Those two systems differ only in how the truth-values are read off these statuses. The degree supervaluationist might assign a truth-value $1/2$ in this case—that is to say, ideally, one would have credence $1/2$ in the borderline proposition. The supervaluationist proper, on the other hand, assigns truth value 0 to the same status, saying one should utterly reject such borderline propositions. Similar differences at the level of truth-value projection discriminate other pairs of theories above (e.g. Kleene vs. LP; Fuzzy vs. fuzzy gaps).²⁶ Thus, even for those who *agree* on a certain kind of non-classical style of semantics, there may well be controversy over the *doxastic role* that the projection onto truth values brings with it. Sometimes this is acknowledged in the literature (the difference between interpreting the ‘third truth status’ as ‘both true’ or ‘neither true nor false’ that differentiates Kleene gaps from LP gluts is familiar, for example). But sometimes it is less prominent. Supervaluationism

²⁵Contrast the generalized dutch book arguments in Paris (2001), where the pivotal role of truth values is *pragmatic*—as specifying what proportion of the prize for a bet on p is obtained in a situation where p has the truth value in question. It’s natural to take the two characterizations to coincide, but they are in principle separable.

²⁶Indeed, ignoring the conditional, we have three projections from essentially the same three-valued statuses: Kleene (which, aside from the conditional, is the same modulo notation as the 3-valued fuzzy gap theory); LP, and the 3-valued Fuzzy system. The ‘third status’ in the respective cases are projected to 0, 1 and $1/2$, respectively.

is a good illustration: it's far from clear what sort of doxastic role, or projection onto truth-values, self-described 'supervaluationists' intend.²⁷ So inevitably, some of the positions I associate with particular systems above will be contentious. Just to illustrate this: I've suggested a 'gappy' interpretation of intuitionism, so that (under the natural interpretation of truth values 1 and 0) some sentences will be neither true nor false. But this interpretation of intuitionism is one deprecated by some prominent intuitionists.²⁸

There's a second role that truth-value projections play. For each kind k of non-classical semantics, we can define a logic \vdash_k . Recall that the status of a proposition $p(w)$, and thus its truth value $\llbracket p \rrbracket_w$, are relative to an interpretation function I . We can read the above characterizations of non-classical semantics as giving constraints, in terms of \neg , \wedge and \rightarrow , on what I should be like—and interpretations

²⁷To some extent the discussion of whether 'truth' can be identified with 'supertruth' might track this distinction.

²⁸Cf. Wright (2003). One worry motivating the move away from gaps is that, combined with natural disquotational principles for truth, lead to contradiction—see Wright (1992). This tension with disquotational effects all the 'gappy' views taken above.

One general line on this (in the spirit of McGee & McLaughlin (1994) and Field (2003)) is to distinguish 'truth values'—which earn their keep philosophically by wider theoretical role e.g. by norming credences—from the behaviour of truth-predicate that obeys the disquotational principles. If that's right, then the 'natural identification' of truth value 1 with truth simpliciter is brought into doubt, and the label 'truth value gaps' is less appropriate.

But I think it's clear that this wouldn't satisfy Wright, who would I think disagree even with the *doxastic* role envisaged for intuitionism here. In the case of vagueness, for example, a borderline case is not supposed to prompt utter *rejection* (as it might seem here) but rather *agnosticism*—a principled kind of 'suspension of judgement' $A \vee \neg A$. See the discussion of quandary in Wright (2001).

In connection with this, one point of resistance to our entire picture is the presupposition that there is an 'intended' interpretation to be had, at which the sentences we utter or thoughts we think can be said to have 'statuses' or 'truth values' of the kind sketched above. Field is explicit in denying a role for 'intended interpretations' in this sense, and intuitionists like Wright are typically rather sceptical of the kind of substantive role for semantics we here suppose.

that meet these constraints will be called *k-admissible*.

k -consequence is then given by the following: $A \vdash B$ iff for each w and k -admissible I , $(\llbracket A \rrbracket_w \leq \llbracket B \rrbracket_w)$. It's easy to check that by construction, the above characterizations give us familiar logical systems:

Classical Classical logic.

Kleene gaps Strong Kleene logic.

LP gluts The 'logic of paradox' LP.

Fuzzy Lukiesciewicz degree-preservation logic

Finite Fuzzy Lukiesciewicz finitely valued logics under degree-preservation.

Fuzzy gaps Lukiesciewicz 1-preservation logic (either finitely valued or continuum valued).

Supervaluations 'Global' supervaluational logic (when restricted to boolean combinations of atomic propositions, this extensionally coincides with classical logic).²⁹

Degree supervaluations Supervaluational degree preservation logic (when restricted to boolean combinations of atomic propositions, this extensionally coincides with classical logic).³⁰

Intuitionistic Intuitionistic logic

In the cases where the truth values are $\{1, 0\}$, from the point of view of fixing the logic, one can think of $\llbracket \bullet \rrbracket$ as simply picking out which

²⁹On the distinction between global and local consequence, and various other possible definitions, see Varzi (2007) and the references therein.

³⁰The relationships between different ways of characterizing consequence in this system are not much discussed—but the various options extensionally coincide for a basic boolean language. For discussion of the more general case, see (reference suppressed).

statuses are ‘designated’ (i.e. the inverse image of 1). ‘No drop in truth value’ becomes, in effect, the requirement that we never have a case where premises have a designated truth status, and the conclusion has an undesignated status. The general characterization of logic in terms of ‘no drop in truth value’ generalizes this to cases where the values are more fine-grained than the simple designated/undesignated distinction allows.

It’s one thing to define a logic for a given setting (one among many possible characterizations), quite another to justify that definition as non-arbitrary and theoretically interesting. As flagged earlier, we will see that the characterizations just given do have this status, once we identify the role they play in constraining coherent credence.

4 Gradational accuracy

We have presented several rival views on the distribution of truth values across propositions. What we’ll now argue is that Joyce’s arguments give precise constraints on rational credences relative to each of the views sketched above. More precisely, Joyce’s *formal* arguments generalize almost immediately to the non-classical case. In his terms, we can show for a certain class of ‘coherent’ credences, that any credence function that is *incoherent* is ‘dominated’ by some coherent credence d —i.e. that d is inevitably more accurate than c , no matter how the actual truth values turn out.

Recall that the basic idea is to evaluate each belief state, at a world, in terms of how close to the truth its credences are overall (with the presumption that, from the epistemic point of view, the closer one is to the actual truth values, the better). To formulate this idea of an accuracy norm in a tractable way, Joyce uses the notion of an ‘inaccuracy score’. First, some setup. A *credence* function represents an assignment of ‘degrees of belief’ to each proposition in Ω . We won’t initially assume that credences have any particular structure—in the general case they’re simply mappings from propositions to positive real num-

bers. The set of all credence functions we call \mathcal{B} . For each world w , there is a very special credence function c_w —one where the credences exactly match the truth values. If the aim of having credences is to match the truth value, then this is the unique maximally accurate credence function, at w . With this in place, we introduce the accuracy score. For each world w and credence function c , $I(c, w)$ measures the inaccuracy of c at w . The injunction to minimize inaccuracy is then understood as the injunction to minimize this quantity. A dominated credence c is one where there is a c' such that $I(c, w) > I(c', w)$ for every $w \in W$.

One way of developing this view is to specify directly what I is to be. A famous option is the *Brier score*:

$$I(c, w) = (1/|\Omega|) \sum_{p \in \Omega} |c(p) - \llbracket p \rrbracket_w|^2$$

(As before, $\llbracket p \rrbracket_w$ denotes the truth value of p at w). Notice that this makes perfect sense even when the truth-values concerned are non-classical.

An alternative, favoured by Joyce, is to lay down certain characteristics any ‘reasonable’ scoring function should satisfy. One then derives results for *any* scoring function meeting the axioms.³¹

³¹Joyce’s axioms are:

1. Structure: For each $w \in V$, $I(b, w)$ is a non-negative, continuous function of b that goes to infinity in the limit as $b(X)$ goes to infinity for any $X \in \Omega$.
2. Extensionality: At each possible world w , $I(b, w)$ is a function of nothing other than the truth-values that w assigns to propositions in Ω , and the degrees of confidence that b assigns these propositions.
3. Dominance: If $b(Y) = c(Y)$ for every Y in Ω other than X , then $I(b, w) > I(c, w)$ iff $|\llbracket X \rrbracket_w - b(X)| > |\llbracket X \rrbracket_w - c(X)|$.
4. Normality: If $|\llbracket X \rrbracket_w - b(X)| = |\llbracket X \rrbracket_u - c(X)|$ for all X in Ω , then $I(b, w) = I(c, u)$.
5. Weak convexity: Let $m = (\frac{1}{2}b + \frac{1}{2}c)$ (i.e. the midpoint of the line-segment between b and c). If $I(b, w) = I(c, w)$, then it will always be the cases that $I(b, w) \geq I(m, w)$, with equality iff $b = c$.

Joyce’s central result is that (for any scoring function meeting a these axioms—which includes in particular the Brier score), any credence assignment that is ‘incoherent’ in the sense of violating the probability axioms, is accuracy-dominated by some coherent credence function. The result is proved under the assumption that the truth-value assignments are classical. But we’ll see that Joyce’s proof is more general than this.

We noted above that to each world there corresponds a ‘perfect match’ credence function c_w , such that $c_w(p) = \llbracket p \rrbracket_w$. Assuming classicism, the probability functions are just weighted averages of these perfect credences (i.e. if $c_1 \dots c_n$ are the perfect credence functions, for each probability function p , there is $\lambda_1 \dots \lambda_n$ such that (i) $\sum_i \lambda_i = 1$, and (ii) for each proposition X , $p(X) = \lambda_1 c_1(X) + \dots + \lambda_n c_n(X)$). That is, the set of ‘coherent’ credence functions V^+ in Joyce’s sense is the ‘convex hull’ of the set of perfect credence functions generated by the set of worlds V .

The generalization of Joyce’s result is essentially that we can let V be the set of perfect credence functions relative to any of the kinds of classical or non-classical truth-value assignments described earlier—and define from this a ‘coherent’ set of credence functions V^+ formed by taking weighted averages on V . We can then prove:

For any $c \in B - V^+$, there is some $d \in V^+$, such that for every $w \in V$, $I(c, w) > I(d, w)$.

In the particular case where V^+ is the convex hull of a set generated by truth-value assignments, the domination result is an immediate corollary of this. To the extent that Joyce, in the classical setting, succeeds in making a case that credences that violate the axioms of probability are irrational because accuracy-dominated, we

can argue that credences that are incoherent by *non-classical* lights are irrational.


5 The generalized proof

This section will give a flavour of the argument that Joyce uses to prove the accuracy domination theorem, without provided the proof itself. We’re working with a space of belief states B , some among which ‘perfectly match’ truth value distributions (call the set of those that do, V). The belief states can be thought of as N -tuples of real numbers, where N is the number of propositions we are working with; so belief space can be modelled by \mathbb{R}^N . The *coherent* credences are just the weighted averages or ‘mixtures’ of elements of V . We call this set V^+ . When V contains only classical truth-functions, the coherent credences are exactly those that satisfy the axioms of classical probability.


The argument for accuracy-domination would be straightforward if ‘inaccuracy’ behaved exactly like the Euclidean distance relation on \mathbb{R}^N —if the ‘inaccuracy’ of b at world w is just given by the Euclidean distance between b and the belief state c_w that exactly matches the truth values at w . The recipe for finding a point that accuracy-dominates b if b is not in V^+ is simply to find the *closest* point to b that *is* in V^+ . To argue for accuracy-domination, we reason geometrically. Suppose that c did not accuracy-dominate b , so that cw is no longer than bw in the following picture:³²

6. Symmetry: If $I(b, w) = I(c, w)$, then for any $\lambda \in [0, 1]$ one has $I(\lambda b + (1 - \lambda)c, w) = I(\lambda c + (1 - \lambda)b, w)$.

³²I’ve drawn the angle bwc as acute. If it were non-acute, then w would be closer to b than c is, which contradicts the choice of c .



We could then drop a line from c perpendicular to cw . The point x where they meet will lie on the line wc , and hence will be a convex combination of w and c , which means (since they are in V^+ , and that set is convex closed) that x is in V^+ :




But the hypotenuse of a right-angled triangle is always the longest side. That means that the distance from x to b is shorter than that between b and c . But this contradicts the choice of c ! For c was supposed to be the closest element in V^+ to b , and x is closer, and in V^+ . By reductio, we conclude that c accuracy dominates b .

Now, all this reasoning presupposes that inaccuracy is measured by the Euclidean distance between a credence and the world (or some order-preserving function of it). Taking the inaccuracy measure to be the *Brier Score* does exactly this. But it's far from obvious that the reasoning can be replicated more generally. Indeed, if all we're given is an inaccuracy measure, relating worlds to credences, there's no guarantee that there's any 'distance' measure relating arbitrary

pairs of belief states even to run the argument with. But furthermore, the informal reasoning above made use of angles (e.g. 'right angled triangle') and even if we found a relation with enough structure to deserve the name 'distance', it's not immediate that talk of angles would be thereby legitimated.

Joyce therefore faced two challenges: to characterize a notion of distance in terms of inaccuracy, and then to use it to rerun the argument. On the first point, Joyce's proposal is to set $D(a,b) := I(w + (a - b), w)$. One can think of $w + (a - b)$ as a point in belief space which stands to w in a way that parallels the way that a stands to b , as pictured below:



The axioms on credence that Joyce imposed turn out to be sufficient to show that the above returns the same distance measure no matter which w we use in its construction; and further, impose *enough* of the behaviour of a distance function to allow Joyce to reconstruct the pattern of argument sketched above. The detailed proof can be found in Joyce's paper: the above geometrical pictures can usefully serve as a heuristic in understanding the various steps.

What is important for our purposes is that to reconstruct this proof, what matters is just that the distance-structure defined is

well-behaved. That V^+ is the convex closure of V is crucial to the proof, but *no* assumptions are necessary about which points are or aren't included in V . In particular, it doesn't matter whether they're classical truth-value distributions (as Joyce assumes) or something else entirely. The practical upshot is this. Begin with V any kind of 'truth-value distributions' you like—degree theoretic, supervaluational, gappy, whatever. Define a notion of 'coherent credence' for that kind of truth-distribution as just sketched. Then Joyce's argument will immediately give us that any incoherent credence (in the new sense) will be accuracy-dominated by a coherent credence (again in the new sense). So we have a very general characterization of 'admissible' credences, relative to all sorts of non-classical approaches to truth-values.³³

Of course, the assumptions about the nature of I that Joyce makes in order for this to go through can be (and have been) challenged. But I'm aiming here for parity: the claim that *if* accuracy considerations make a convincing case for probabilism relative to classical semantics, they *also* give a recipe for the analogues of probabilism just described, relative to non-classical semantics.

It's worth noting that this neutrality is *not* a general characteristic of all arguments for accuracy-domination. For example, in later work, Joyce (2009) gives a very different accuracy-domination argument for probabilism. In outline, he argues: (i) for an arbitrary partition Δ , there is an accuracy measure I_Δ that is sensitive only to the degrees of belief in proposition drawn from that partition. (ii) it is shown (given axiomatic constraints on partition-relative accuracy) that if these credences don't add to 1, then accuracy-domination results; (iii) one argues that if a belief state is not a probability, then there will be some partition the credences of the cells of which do not add to 1. This shows that non-probabilistic belief states will be 'accuracy-dominated'

³³Compare the use of Brier-score domination (with a very different philosophical spin) to constrain expectations of random variables in De Finetti (1974).

on some partition.³⁴

It turns out that the axiomatic constraints required by (ii) are far simpler, and perhaps more appealing, than Joyce's originals. However, the evaluative significance of accuracy-domination *relative to a partition* is far less transparent than the original idea of overall accuracy. But however this works out, what's crucial for our purposes is to note that the *structure* of the argument makes appeal to logical features of the propositions it works with—we need the notion of propositions forming a partition, for example. Such notions do not automatically transfer to non-classical settings, but must be thought through anew. It turns out that there is an understanding of 'partition' that suffices for (ii) to go through—a set of propositions is a partition in this sense if at every world, the sum of the truth values of members of that set is 1. But it's not obvious that there's a *general* argument available that coherence 'on each partition' in this sense suffices for overall coherence. Joyce's later argument puts into sharp relief the *neutrality* of his earlier argument.

6 Domination immunity

If we were working with the Brier score, then we can prove the converse to the above accuracy-domination theorem: any *coherent* credence will *not* be accuracy-dominated. This holds whether we work in a classical or non-classical setting.³⁵ But it's not at all clear whether Joyce's axioms suffice to show that coherent credences are domination-immune in this way. I know of no proof one way or the other on this front. Where this issue has been raised in the literature, a typical

³⁴It's not obvious from the text that this is the shape of the overall argument; but in p.c. Joyce suggested this construal.

³⁵This follows from the analogous result for synchronic dutch books in the generalized setting (Paris, 2001), together with the observation that every point that accuracy-dominates a belief state can be used to construct a dutch book for that state (Author, suppressed).

move is to consider supplementations to Joyce's axioms.³⁶

Perhaps the supplementations can be justified; and perhaps Joyce's axioms unsupplemented suffice for the converse result. This would certainly be a welcome result. But give the present state of the literature, it's worth assaying the philosophical significance if the converse result *failed*. We should therefore distinguish the set C of convex combinations of truth-value assignments from the set D of domination-immune belief states. Joyce's result shows that $D \subseteq C$. A converse would show that $D = C$. We know also that D is at least non-empty (the worlds themselves are members). What follows if $D \neq C$?

Suppose pro tem that accuracy-domination is acknowledged as sufficient for the lack of ideal rationality, so that the upshot in such a scenario is that some probabilistic belief states, as well as some non-probabilistic ones, are bad because accuracy-dominated. It's not *obvious* why this is a bad result. Probabilistic constraints would retain their normative force—a belief state that violates them would be bad. One could not conclude simply from the fact one satisfied probabilistic norms that one's belief state was rational. But this wouldn't follow in any case, unless probabilistic constraints were the *only* rationality constraint in force. Even subjective Bayesians think that there are at least diachronic consistency constraints; and most others would be open to further constraints—respect for the known chances, for

³⁶Joyce (2009) argues that it is reasonable simply to require of an accuracy-measure that it makes all credence functions domination-immune (his argument is that a credence function shouldn't count as flawed if it matches a possible distribution of objective chances; and each probability function describes such a possible distribution. Pettigrew (Manuscript) adds the principle of dominated-compromise respect. This says that if b, b', c, c' are such that $\forall w I(b, w) \leq I(c, w)$ and $I(b', w) \leq I(c', w)$ then $I(\lambda b + (1 - \lambda)b', w) \leq I(\lambda c + (1 - \lambda)c')$. Informally: if b dominates c and b' dominates c , then a mixture between c and c' will be dominated by the corresponding mixture between c and c' . This additional axiom says that the set of domination-immune belief states is closed under convex combination, and since the other axioms already guarantee that at least the worlds are domination immune, we have that the convex combinations thereof (the probabilities) are domination-immune. Pettigrew's axiom does the job

example. So it is not clear, at first, what the downside is meant to be.

I see three potential complaints. The first complaint is that of *overdemandingness*—for all we've said, almost all belief states could be accuracy-dominated, in which case the claimed irrationality would be implausibly strong. The second is the *undermining* constraint—that the lack of a converse result would undermine the original case that accuracy-dominated credences are bad. The third complaint is that this account would be objectionably *revisionary of wider theory*—regarding some probabilistic belief states as irrational in this sense would undermine well-entrenched resources.

The first worry is that there might be ways of measuring accuracy, compatible with Joyce's axioms, that make almost all belief states accuracy-dominated. Without the 'safety result' that probabilistic credences are domination-immune, there's no blanket guarantee that a given sensible-looking belief state won't have this 'flaw'. But rather than a discovery, such a prospect would count as a *reductio* of the approach.

The crux of the first worry concerns the status of the axioms. One construal is that they merely set down some constraints that we are justified in thinking that the accuracy measure satisfies. They need not be taken to be all the information we have. Whether an accuracy-measure is acceptable may be determined by whether it fits general principles such as the Joyce axioms, but also whether it returns the right verdicts in paradigm cases. Maybe the One True accuracy measure (or the admissible candidate accuracy-measures) should be taken to be the maximally simple or natural measure(s) meeting the previous two desiderata. In virtue of the first constraint, we know that improbable belief states will be accuracy-dominated; in virtue of the second, we know that accuracy-domination won't return the wrong results in paradigm cases. The worry might be that the three criteria just mentioned would conflict, so that every candidate accuracy-measure is ruled out. But the existence of a Brier score that fits the constraints, doesn't exclude paradigm probabilistic belief states, and

is maximally simple, shows that we don't get into that unfortunate situation.

The second worry is that even if D is large enough to be plausible, the lack of a converse to Joyce's result undermines the argument that not being in D is sufficient for lack of (ideal) rationality. Suppose b is an improbableistic belief state. Then by Joyce's result, it is dominated by some probabilistic c . But there's no guarantee that c is in D —it might in turn be dominated by some (probabilistic or improbableistic) c' , and so on ad infinitum. (Hájek, 2008, cf.)

There are weaker and stronger ways of understanding the complaint. Some regard this as undercutting the charge that accuracy-domination is a flaw in b . The defence responds that the charge is not undercut—having belief state b , as opposed to the necessarily more accurate c , is still a bad thing; that c is likewise dominated by some other belief state just shows the considerations iterate.

A better version of this complaint would be not that accuracy-domination becomes ok when it's embedded within an infinite chain; but that (even ideal) agents may be stuck in a bad situation with no responsible way to extricate themselves. Perhaps some others *don't* face this complaint—but councils of perfection don't matter too much to someone who hasn't a responsible route to achieve it. Given a converse result and an improbableistic starting belief state, we are able to point to a probabilistic credence that not only accuracy-dominates the starting point, but also itself is domination-immune. So moving to that state removes the flaw, and is also a justifiable move in itself, insofar as it improves accuracy. But in the envisaged situation, for all we've said so far, the only way to get from the starting point to a domination-immune belief state may be to arbitrarily shift your beliefs. But this is not an option for a responsible epistemic agent. The significance of a converse to Joyce's result is that it shows that there's at least one belief state in D which is more accurate than the starting point, and so gives our agent a responsible route into D .

So construed, the complaint has force; but it can be addressed.

Here is one path we could offer an agent with an incoherent starting point. She should move first from her incoherent b to coherent c , and then evaluate coherent belief states by their *expected inaccuracy* by c 's lights. Since we're working with coherent belief states, this is a well-defined notion. Expected inaccuracy is a continuous function of the points in belief space, and so achieves a minimum within the closed convex set of coherent credences. Shifting from c to that state which minimizes expected inaccuracy by c 's own lights is a reasonable way to update. The resulting belief state d cannot be accuracy-dominated, or it would not minimize expected inaccuracy. Proof: if d is dominated, it will be dominated by some coherent credence by Joyce's result and the transitivity of dominance. But this can't happen, since if d were dominated by e , e would be uniformly more accurate than d , hence *on average* more accurate than d , so the latter would not minimize expected inaccuracy, contrary to its construction. So this gives us a reasonable route from an arbitrary starting point outside of D into D itself.³⁷

The final complaint against the lack of a converse is that without a general guarantee that (generalized) probabilities are immune from accuracy domination, there's no guarantee that working with

³⁷One might worry about a revenge puzzle. The belief state d , we've said, can't be accuracy-dominated. But it might not be 'self-confident'—it might fail to minimize expected inaccuracy by its own lights. In which case, a second regress might begin. However (i) this isn't a new issue—even with a converse to Joyce, there's no obvious guarantee that the above situation won't arise; (ii) the regress is only damaging if it's a flaw in a belief state that it fails to minimize expected inaccuracy by its own lights. It's not clear to me that this is the case—the above argument only required that the shift from one belief state to another via expected accuracy minimization was permissible, not that it was mandatory. It's worth noting that a belief state might not minimize expected inaccuracy by its own lights, but still be 'stable'—a person might think that e is the best belief state to have, and falsely believe that they are in e . This may violate 'special reflection' (van Fraassen, 1984); but then it is the status of the latter principle of doxastic transparency we should be discussing, not expected inaccuracy minimization. Compare (Leitgeb & Pettigrew, 2010a).

the set of ‘rational’ belief states D will fit with other, well-entrenched resources—and so would count as objectionably revisionary. For example, it may be that updating procedures (like conditionalization) won’t take us from an undominated point to a dominated point. So a generalized version of standard Bayesian techniques can’t be considered *safe* without a converse to Joyce’s argument.

One response would be to view this as a variant of the over-demandingness worry, and to respond in the same way. But a more direct response is available. First, if there are substantive rationality constraints on belief beyond the bare probabilistic/logical ones, there’s already a danger that conditionalization will move us to an ‘irrational’ spot in belief space.³⁸

A plausible reaction to these cases is to reach for a more general characterization of updating—one that agrees with conditionalization under simplifying assumptions, but is compatible with more sophisticated constraints being in place. Objectivist Bayesianisms have long recommended that we update by advising us to move to that probability that maximizes informational cross-entropy, while meeting evidential and rational constraints.³⁹ If the constraints are simply that a certain proposition be assigned credence 1 within a probabilistic belief state, then this gives us conditionalization. But there’s nothing in principle that precludes adding ‘dominance immunity’ (or special reflection, or a conditional probability-probability conditional principle) to the constraints under which one maximizes entropy. So even if dominance-immunity as a rationality constraint may be *somewhat*

³⁸Moore-paradoxical propositions may already be a case of this (van Fraassen, 1984). Another interesting potential case is updating on conditionals. One reaction to the diachronic triviality results that were proved following Lewis is that they show that the requirement that the credence in a conditional matches the corresponding conditional credence needs to be imposed as an exogenous rationality constraint, and is violated if we update by conditionalization.

³⁹See Williamson (2010) for recent defence of this form of updating; and Leitgeb & Pettigrew (2010b) for discussion of what happens if we use the Brier score and the induced measure across belief states to handle updating under a constraint.

revisionary in the absence of a converse to Joyce’s result, there are extant generalizations that take it in their stride.

In summary, I see the principle significance of the Joyce result to be that probabilistic constraints have normative force; and I think there are good grounds for rejecting the Hajekian claim that this force is undermined if we lack a converse. What is certainly true is that that it would be *interesting* to show the converse of Joyce’s result; or failing that, to limn the boundaries of the domination-immune belief states. This is a matter for future research. But the philosophical significance of the results so far established does not wait on it.

7 Axiomatizing coherent credences

We turn next to the question of whether we can find useful ways of capturing the constraints given above. A model for this is the role that probability axioms (and theorems) play in the classical case. This captures general patterns which a belief state must satisfy in order to count as coherent—giving more ‘local’ versions of the general requirement that one’s credences be representable as mixtures of truth value assignments. One thing that is highly significant here philosophically is the role that logical *entailment* plays in formulating these coherence constraints. But the axiomatization translates a constraint formulated in terms of *truth values* into principles that talk about how logic should inform your credence distributions. An argument for coherence in this sense, together with a formulation in terms of a particular notion of logic, allows us to *derive* a normative role for logic.

This is interesting enough in the classical case. But it’s particularly significant in the non-classical case, where it’s often not clear what the principled generalization of consequence should be (or even if there’s any clear content to the question of which, among several inequivalent definitions, is the true generalization of consequence). If there is one characterization in a non-classical setting which preserves the normative role associated with classical logic, it will have a good

claim to ‘play the logic role’.

We are going to look at two kinds of results. The first is a general observation: there’s a particular uniform characterization of consequence in the non-classical setting that’s well-suited to play the role mentioned. This is ‘no drop’ consequence: a guarantee that there is no drop in truth value over a valid argument. It’s fairly easy to show in a wide variety of settings, coherent credences must interact with this notion of consequence as classically coherent credences do with classical validity. The second is a more detailed set of results: to use such a characterization to formulate axiomatizations of non-classically coherent credences that are *necessary and sufficient* for credence. The significant mathematical work here is in showing that a given set of axioms are complete, in the sense that satisfying them is *sufficient* for being a convex combination of truth-values of the relevant kind. Paris (2001) draws together a number of such results in the paper cited; they are critically discussed below.

Paris’s characterization build on the following axiomatization of classical probabilities, where \vdash_k is understood as classical consequence:

$$\begin{aligned} (\mathcal{P}1) \quad & \vdash_k A \quad \Rightarrow \quad b(A) = 1 \\ & A \vdash_k \quad \Rightarrow \quad b(A) = 0 \\ (\mathcal{P}2) \quad & A \vdash_k B \quad \Rightarrow \quad b(A) \leq b(B) \\ (\mathcal{P}3) \quad & b(A \wedge B) + b(A \vee B) = b(A) + b(B) \end{aligned}$$

It’s well known that meeting the following three constraints is necessary and sufficient for being a convex combination of classical truth values. The central result in Paris’s paper is that the same three axioms characterize the convex combinations of truth values so long as (i) truth values are taken from $\{0, 1\}$; (ii) $A \vdash_k B$ is given the ‘no drop’ characterization mentioned earlier;⁴⁰ and (iii) the following are satisfied:

⁴⁰Paris formulates this in terms of the preservation of value 1 across the sequent, but given (i), these are equivalent.

$$\begin{aligned} (\mathcal{T}2) \quad & V(A) = 1 \wedge V(B) = 1 \iff V(A \wedge B) = 1 \\ (\mathcal{T}3) \quad & V(A) = 0 \wedge V(B) = 0 \iff V(A \vee B) = 0. \end{aligned}$$

This result applies directly to many of the settings set out earlier: Classical, Kleene Gaps, LP gluts, Fuzzy Gaps, and Intuitionistic frameworks all satisfy Paris’s axioms.⁴¹

Drawing on the work of Gerla (2000) and Di Nola *et al.* (1999), Paris argues that a similar result holds for finite fuzzy (Łukasiewicz) setting—and Mundici (2006) later extended this to the continuum valued fuzzy setting. As in the classical case, ($\mathcal{P}2$) turns out to be redundant. Again, the crucial maneuver is the change in interpretation of \vdash_k and here it is crucial that we understand this as the ‘no drop’ logic given earlier—the 1-preservation logics that are often referred to as ‘Łukasiewicz’ logics *will not do*.⁴²

The gap supervenational and degree supervenational settings remain to be covered. The latter is handled quite straightforwardly *for cases where the language only contains standard propositional oper-*

⁴¹For the intuitionistic case, compare Weatherson (2003). Paris reports the general result as a corollary of a theorem of Choquet (1953)

⁴²In his paper, Paris states the result by giving just the first and third axioms (his 1 and 3), and states that the turnstile picks out ‘Łukasiewicz logic’, which might be thought to refer to the 1-preservation understanding. But consider a sentence of the form $A \wedge \neg A$. In a 3-valued fuzzy setting, this can take the values $\frac{1}{2}$ or 0, but never 1. This makes it inconsistent in the 1-preservation sense, and so the second half of 1 would claim that any convex combination of truth value should assign it zero. But clearly this is not correct. And in fact, in the Gerla paper that Paris cites, the condition that plays this role is that A take value 0 on every truth value distribution—which is the condition for being a no-drop inconsistency, not a 1-preservation inconsistency. If one wished to state the conditions in terms of the 1-preservation logic, it is possible to do so—for the second half of $\mathcal{P}1$ could be replaced by the condition that $b(\neg A) = 0$ if A is a 1-tautology. However, $\mathcal{P}2$ could not then be retained, and so it’s pretty clear that it’s the no drop, rather than 1-preservation logic that generalizes the standard normative role. (Note that by the earlier result, $\mathcal{P}1$ –3 formulated in terms of 1-preservation directly are necessary and sufficient characterizations of coherent credences appropriate to a fuzzy *gap* setting.)

ators. Recall that every degree supervenational truth-value distribution is induced by a measure over the set of sharpenings (classical truth-value assignments). We could equally use that measure to construct a classical probability assignment over the same sentences. Convex combinations of such truth value distributions will match convex combinations of the associated probabilities—which returns another probability—so every convex combination of degree supervenational truth values will match some classical probability. Conversely, every classical probability is representable by a measure over the space of classical truth value distributions which we can use to induce a matching degree-supervenational truth-value assignment. So classical probability axioms (formulated in terms of classical logic) will characterize the degree-supervenationally coherent belief states. The no drop logic over this language exactly coincides with classical logic, so this generalizes the pattern we have seen so far. The caveat about the expressive resources of the language is crucial, however—as we’ll see shortly.

The one case the pattern breaks down is with the gap supervenational truth value distribution. This case is covered by a theorem that Paris gives drawing on the work of Shafer (1976) and Jaffray (1989). The interpretation is rather different—Paris, like Shafer, gives an *epistemic* gloss on the value 1 as ‘known true’ and 0 as ‘not known true’. But this doesn’t matter for the technical result. For the propositional language under consideration, the results show that convex combinations of such truth values state are exactly the *Dempster-Shafer belief functions*. These may be axiomatized thus:

$$\begin{aligned}
 (\mathcal{DS}1) \quad & \vdash A \quad \Rightarrow \quad b(A) = 1 \\
 & A \vdash \quad \Rightarrow \quad b(A) = 0 \\
 (\mathcal{DS}2) \quad & A \vdash B \quad \Rightarrow \quad b(A) \leq b(B) \\
 (\mathcal{DS}3) \quad & b(\bigvee_{i=1}^m A_i) \geq \sum_S (-1)^{|S|-1} b(\bigwedge_{i \in S} A_i)
 \end{aligned}$$

(where S ranges over non-empty subset of $\{1, \dots, m\}$). Paris’s initial formulation is slightly different, and uses classical logic (p.7), but as

he notes this is extensionally equivalent to current version using the ‘no drop’ logic over the ‘supervenational’ truth values (p.10).

These later results are proved *for specific languages* rather than whole classes of languages. We do *not* have a guarantee that adding expressive resources preserves the result. For the gap-supervenational settings, this is absolutely crucial. In discussions of vagueness, it’s standard to study a language containing a ‘determinately’ operator D —with Dp true on a sharpening iff p is true on all sharpenings. And it’s well known that adding this resources introduces non-classical behaviour into the ‘no drop’ gap supervenational logic (‘global supervenational logic’). For example, we have the following pair, showing that the classical metarule of reductio fails in this setting:⁴³

$$\begin{aligned}
 & p \wedge \neg Dp \vdash \\
 & \not\vdash \neg(p \wedge \neg Dp)
 \end{aligned}$$

The proposition $p \wedge \neg Dp$ is always value 0; but its negation can be 0 too. On every convex combination of supervenational truth values c , $c(p \wedge \neg Dp) = 0$ —so a ‘complete’ set of axioms for coherence should enforce this. Likewise, in the degree supervenational setting, $p \wedge D_{0.5}p$ (where the latter conjunct is true on a sharpening when p is true on half the admissible sharpenings) is always either value 0 or 0.5; and this constraints coherent credences in the proposition.

Axiomatizations of coherent credences over a set of propositions that don’t include determinately operators may fail when we take them into account. For example, in the supervenational settings, we could (and Paris originally does) state the axioms using classical logic. But for one thing, it’s not clear how this should extend to a language with distinctively non-classical vocabulary like D . And for another, it seems plausible that to capture this behaviour, we need the power of the no-drop logic, and in particular, the fact that it makes the conjunction above inconsistent. The formulation given above is

⁴³Cf. (Williamson, 1994, ch.5).

designed to preserve as much power as possible in these expressively richer settings; but even so we can't rely on the original completeness result.⁴⁴ The completeness results are therefore *fragile*.

Matters are even worse in the degree-supervaluational setting. In (Author, suppressed) I argue that to fully capture the constraints on degree-supervaluational credences we need to introduce extra resources—a generalization of $\mathcal{P}2$ to cover a multi-premise case, and perhaps even a range of degreed consequence relations. Finding a formulation for which a completeness result is even plausible is a hard, open-ended project.

But even if we don't have a *completeness* result for an axiomatization, it will still be interesting if the characteristic logical constraints on coherent credence we've identified $\mathcal{P}1$ – 3 remain *constraints* under expansions of the language. That would mean that their normative force is not conditional on expressive limitations. Perhaps there may be credences meeting the axioms which aren't converse combinations of truth values. Given our earlier discussion of the lack of a converse to Joyce's theorem, perhaps there are converse combinations of truth values that are accuracy-dominated. Even if both are realized, we can still hope to show that violations of characteristic logical constraints on credence *guarantee* irrationality in the form of accuracy-domination. Tightening the connections in the form of completeness results (for the first lacuna) and a converse to Joyce's results (in the second) would be interesting and significant. But we could establish the normative force of logic on credence without it; and this where much of the philosophical interest lies.

The argument that the axioms are binding when stated using the appropriate no-drop logic tends to be straightforward. We'll briefly run through what's involved. By earlier results, if c is a coherent

⁴⁴A further issue concerns the shape of the axioms. Paris's original formulation, drawn from Shafer and Jaffray, formulates constraints on credences only appealing to tautologies, rather than valid sequents. But as we've seen above, once D is in the language, the classical equivalence between inconsistency and tautology of the negation breaks down, allowing further inequivalences between formulations.

credence, and t_i are the truth value assignments, $c = \sum_i \lambda_i t_i$, for suitable choices of $\lambda_i \geq 0$ such that $\sum_i \lambda_i = 1$. Note that for all the systems we're considering, if $A \vdash_k$, then for all i , $t_i(A) = 0$. Thus $c(A) = \sum_i \lambda_i t_i(A) = \sum_i \lambda_i \cdot 0 = 0$. Equally, if $\vdash_k A$, then for all i , $t_i(A) = 1$. Thus $c(A) = \sum_i \lambda_i t_i(A) = \sum_i \lambda_i \cdot 1 = \sum_i \lambda_i = 1$. So the second axiom holds. From the characterization of \vdash_k , if $A \vdash_k B$, then for any i $t_i(A) \leq t_i(B)$. It follows (since $\lambda_i \geq 0$) that $c(A) = \sum_i \lambda_i t_i(A) \leq \sum_i \lambda_i t_i(B) = c(B)$. So the analogues of the first two axioms hold for each of the various notions of 'coherent credence' we're dealing with. Note that nothing about the character of the connectives or truth value distributions needs to be assumed for this: it is a quite general guarantee, and remains true no matter what extra expressive resources are added in.

$\mathcal{P}3$ will hold for all those systems where we have:

$$t_i(A) + t_i(B) = t_i(A \wedge B) + t_i(A \vee B).$$

One way to guarantee this is by having conjunction and disjunction satisfy: $t_i(A \wedge B) = \min(t_i(A), t_i(B))$ and $t_i(A \vee B) = \max(t_i(A), t_i(B))$. We need only note that $x + y = \min(x, y) + \max(x, y)$. This holds for all the non-supervaluational settings given earlier. It also holds of the degree-supervaluational settings, where we can read this off the analogous axiom for measures over classical truth-value assignments.

Either way, given this basic result, it follows that

$$\sum_i \lambda_i (t_i(A) + t_i(B)) = \sum_i \lambda_i (t_i(A \wedge B) + t_i(A \vee B))$$

and hence

$$\sum_i \lambda_i t_i(A) + \sum_i \lambda_i t_i(B) = \sum_i \lambda_i t_i(A \wedge B) + \sum_i \lambda_i t_i(A \vee B).$$

as promised. Again, it doesn't matter whether or not *additional* resources are added to the languages, so long as conjunction and disjunction continue to obey the constraint.

What remains to be discussed is the robustness of $\mathcal{DS3}$ in the gap supervalational setting. This follows from the fact that if we have a set of n propositions A_i , for $i \in \Lambda$ then the truth values of the n -ary disjunction satisfies:

$$t_i(\bigvee_{i \in \Lambda} A_i) \geq \sum_{S \subset \Lambda} (-1)^{|S|-1} t_i(\bigwedge_{i \in S} A_i)$$

Since the inequality holds for arbitrary truth-value distributions, it holds for their convex mixtures, just as above, which is just $\mathcal{DS3}$. We relegate the proof of the fact about supervalational truth values to a footnote.⁴⁵

⁴⁵We adapt a standard inclusion-exclusion argument. Let $D := \bigvee_{i \in \Lambda} A_i$, and write $|X|$ for $t_i(X)$. There are two cases to consider: when $|D|$ is 0 and when $|D|$ is 1. In the former, the truth value of each disjunct must equally be zero, and hence the truth value of conjunctions thereof is zero. Hence the LHS of the above is zero, and the inequality holds. So consider the case where $|D|$ is 1. We start from the observation that:

$$(|D| - |A_1|)(|D| - |A_2|) \dots (|D| - |A_n|) \geq 0$$

this multiplies out to give:

$$\sum_{S \subset \Lambda} |D|^{n-|S|} (-1)^{|S|} \prod_{i \in S} |A_i| \geq 0$$

or equivalently (recalling that $|D|=1$):

$$|D| + \sum_{S \subset \Lambda} (-1)^{|S|} \prod_{i \in S} |A_i| \geq 0$$

But the truth value of a conjunction, in the supervalational setting as well as the classical one, is equal to the product of the truth values of its conjuncts (it has truth value 1 iff each conjunct has truth value 1; otherwise it has truth value 0). So we can replace the products in the above equation with the truth value of the relevant conjunction:

$$|D| + \sum_S (-1)^{|S|} \left| \bigwedge_{i \in S} A_i \right| \geq 0$$

QED.

Conclusion

Non-classical truth-value assignments disrupt Joyce's original argument for (classical) probabilism. But as we have seen, a generalized probabilism follows if we grant his premises other than classicism. As a bonus, we get a general rationale for a family of non-classical probabilities, and a reason to be interested in the specific characterization of logical consequence that plays the central role in the local characterization of these probabilities. The accuracy arguments 'bundle together' truth values, logic, and ideal credence in a pleasing way.

One striking feature of the axiomatizations given above is that the third axiom—the one that must be varied between a gap supervalational system and the others—is not framed *purely* in terms of consequence. Rather, it builds in a specific relationship between conjunctions and disjunctions, and their conjuncts/disjuncts. We in effect hard-wire in certain assumptions about the logic of conjunction/disjunction when setting down such an axiom. A natural suggestion is therefore to look for 'pure' axiomatizations of probability theory, which will generalize. I have some suggestions in this direction, but there is far more work to do.⁴⁶

Worries about Joyce's other premises can of course be reraised in

⁴⁶One natural thought is to try to strengthen axiom 2, by appealing to multi-premise consequence. A natural candidate here is Suppes' theorem (writing $\bar{\rho}(X) = 1 - \rho(X)$):

$$\text{If } \Gamma \vdash \psi, \text{ then } \sum_{\phi \in \Gamma} \bar{\rho}(\phi) \geq \bar{\rho}(\psi)$$

which we can show is satisfied by generalized probabilities so long as we adopt the following multi-premise generalization of 'no drop' consequence:

$$\Gamma \vdash \psi \text{ iff } \forall i \forall w (\sum_{\phi \in \Gamma} |\bar{\phi}|_w^i \geq |\bar{\psi}|_w^i)$$

But this on its own won't be sufficient even to replace the ordinary P3. If we introduce a (rather non-standard) multi-conclusion aspect to our consequence relation, we can fix this. The basic idea of this notion of consequence is that there be no more untruth in the conclusions than is already present in the premises. The

the non-classical settings. We can ask about the first step of the three-step argument outlined in the first section. Do we really need to assume that accuracy obeys all the axioms Joyce gives? Can we defend those axioms?⁴⁷ The point raised earlier about various goals for the argument is relevant: circularity constraints on justification have real teeth if we want to argue for the truth of probabilism, but are less constraining if our aim is to explain where the normative force comes from, presupposing probabilism is true. It's worth noting that the move to the non-classical case introduces an extra layer of dialectical complexity, for our main interest may well be in the comparative: *insofar* as we think there is a rationale for probabilism in the classical case, what does that rationale support in the non-classical case? If so, we can in principle use classical probabilism (in the classical case) to motivate constraints on accuracy, and then use that notion of accuracy to argue for the non-classical variants of probabilism set out here.

general characterization of multi-premise, multi-conclusion consequence will be:

$$\Gamma \vdash \Delta, \text{ iff } \forall i, w \left(\sum_{\phi \in \Gamma} |\bar{\phi}|_w^i \geq \sum_{\psi \in \Delta} |\bar{\psi}|_w^i \right)$$

which leads to the following generalization of Suppes:

$$\text{If } \Gamma \vdash \Delta, \text{ then } \sum_{\phi \in \Gamma} \bar{\rho}(\phi) \geq \sum_{\psi \in \Delta} \bar{\rho}(\psi)$$

In the classical setting (and the fuzzy generalizations) we have $A \wedge B, A \vee B \vdash A, B$ and $A, B \vdash A \wedge B, A \vee B$ from which we derive, respectively, $\rho(A \wedge B) + \rho(A \vee B) \leq \rho(A) + \rho(B)$ and $\rho(A \wedge B) + \rho(A \vee B) \geq \rho(A) + \rho(B)$, deriving the original version of $\mathcal{P}3$. The first sequent fails in the generalized supervaluational logic, so we only get the latter inequality, which is a special case of $\mathcal{DS}3$. Inclusion-exclusion results for truth-values will allow us to establish a generalized consequence of this form between the various conjuncts and disjuncts, from which $\mathcal{DS}3$ itself follows.

⁴⁷One sceptic, mentioned earlier, is Maher (2002). For a survey of relevant results, worries about the axioms, and an alternative proof of accuracy-domination, see Joyce (2009). I mentioned earlier some concerns about Joyce's favoured argument for accuracy domination in the 2009 paper; nevertheless, once we're clear about how that is intended to go in the classical case, the project of seeing whether they extend is a natural one.

We can equally ask about the third stage of the three-step model. Does accuracy-domination have the philosophical significance attributed to it? Does it really make the case (whether in the classical or non-classical case) that we *should* have credences that are generalized probabilities?⁴⁸ In other work, I argue that a qualified form of probabilism would follow from accuracy-domination.

It is not only credences to which the accuracy-arguments might be applied. Degrees of evidence for or against given propositions may also be considered. It seems initially plausible to me that if E describes one's degrees of evidences in a body of propositions, then it cannot be accuracy-dominated. If E' is inevitably closer to the truth than E , then surely E' is a better candidate than E for being one's evidence. Accuracy-domination arguments can then be appealed to show that degrees of evidence must be structured probabilistically (in either the classical or generalized sense).⁴⁹

The interest of accuracy-domination results is not hostage to the particular spin that Joyce gave them—interesting though that is. But in each case, the question of how the case generalizes to a non-classical setting arises. This paper provides that generalization.

References

- Broome, J. 2005. Does Rationality Give Us Reasons? 1. *Philosophical issues*, **15**(1), 321–337.
- Choquet, G. 1953. Theory of capacities. *Annales d'institute fourier*, **V**, 131–295.

⁴⁸See Hájek (2008) for discussion. An important point—made to me by NN and MMin particular—is that if the normative claim here is *epistemic* in character, we need to consider how epistemic virtues other than accuracy interact. An accuracy-dominated credence might have holistic epistemic virtues that compensate for its lack-of-accuracy.

⁴⁹Compare Williamson (2000), where a probabilistic notion of evidence is defended.

- De Finetti, Bruno. 1974. *Theory of probability: vol.1*. New York: Wiley.
- Di Nola, A., Georgescu, G., & Lettieri, A. 1999. Conditional states in finite-valued logic. *Pages 161–174 of: Dubois, D, H., Prade, & Klement, E.P. (eds), Fuzzy sets, logics, and reasoning about knowledge*. Kluwer.
- Edgington, Dorothy. 1997. ‘Vagueness by degrees’. *Pages 294–316 of: Keefe, Rosanna, & Smith, Peter (eds), Vagueness: A reader*. Cambridge, MA: MIT Press.
- Field, Hartry H. 1973. ‘Theory change and the indeterminacy of reference’. *Journal of philosophy*, **70**, 462–81. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001) pp. 177–198.
- Field, Hartry H. 2003. ‘Semantic paradoxes and the paradoxes of vagueness’. *Pages 262–311 of: Beall, J. C. (ed), Liars and heaps*. Oxford: Oxford University Press.
- Field, Hartry H. 2008. *Saving truth from paradox*. Oxford: Oxford University Press.
- Fine, Kit. 1975. ‘Vagueness, truth and logic’. *Synthese*, **30**, 265–300. Reprinted with corrections in Keefe and Smith (eds) *Vagueness: A reader* (MIT Press, Cambridge MA: 1997) pp.119–150.
- Gerla, B. 2000. MV-algebras, multiple bets and subjective states. *International journal of approximate reasoning*, **25**(1), 1–13.
- Gibbard, A. 2008. Rational Credence and the Value of Truth. *Oxford studies in epistemology*, 143.
- Hájek, Alan. 2008. ‘Arguments for—or against—probabilism?’. *British journal for the philosophy of science*, **59**.
- Jaffray, J-Y. 1989. Coherent bets under partially resolving uncertainty and belief functions. *Theory and decision*, **26**, 90–105.
- Joyce, James M. 1998. ‘A non-pragmatic vindication of probabilism’. *Philosophy of science*, **65**, 575–603.
- Joyce, James M. 2004. ‘The accuracy of partial beliefs i and ii’. *Lecture notes*.
- Joyce, James M. 2009. ‘Accuracy and coherence: prospects for an alethic epistemology of partial belief’. *Pages 263–297 of: Huber, Franz, & Schmidt-Petri, Christoph (eds), Degrees of belief*. Springer.
- Kamp, J. A. W. 1975. ‘Two theories about adjectives’. *Pages 123–155 of: Keenan, E. (ed), Formal semantics of natural language*. Cambridge: Cambridge University Press. Reprinted in Davis and Gillon (eds) *Semantics: A reader* (Oxford University Press, Oxford, 2004) pp.541–562.
- Keefe, Rosanna. 2000. *Theories of vagueness*. Cambridge: Cambridge University Press.
- Kolodny, N. 2007. How Does Coherence Matter? *Proceedings of the aristotelian society*, **107**, 229.
- Leitgeb, H., & Pettigrew, R. 2010a. An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of science*, **77**(2), 201–235.
- Leitgeb, H., & Pettigrew, R. 2010b. An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of science*, **77**(2), 236–272.
- Lewis, David K. 1970. ‘General semantics’. *Synthese*, **22**, 18–67. Reprinted with postscript in Lewis, *Philosophical Papers I* (Oxford University Press, 1983) 189–229.

- Machina, Kenton F. 1976. 'Truth, belief and vagueness'. *Journal of philosophical logic*, **5**, 47–78. Reprinted in Keefe and Smith (eds) *Vagueness: A reader* (MIT Press, 1997) pp.174-204.
- Maher, Patrick. 2002. 'Joyce's argument for probabilism.'. *Philosophy of science*, **69**, 73–81.
- Maudlin, Tim. 2004. *Truth and paradox: Solving the riddles*. Oxford: Oxford University Press.
- McGee, Vann, & McLaughlin, Brian. 1994. 'Distinctions without a difference'. *Southern journal of philosophy*, **supp XXXII**, 203–251.
- Mundici, D. 2006. Bookmaking over infinite-valued events. *International journal of approximate reasoning*, **43**(3), 223–240.
- Paris, J.B. 2001. A note on the Dutch Book method. *Pages 301–306 of: Proceedings of the second international symposium on imprecise probabilities and their applications, isipta*. Ithaca, NY: Shaker.
- Pettigrew, Richard. Manuscript. Accuracy, chance and the principal principle.
- Priest, G. 2006. *In contradiction: a study of the transconsistent*. Oxford University Press, USA.
- Priest, Graham. 2001. *An introduction to non-classical logics*. Cambridge: Cambridge University Press.
- Shafer, G. 1976. *A mathematical theory of evidence*. Princeton: Princeton University Press.
- Smith, Nick J. J. 2008. *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- van Fraassen, Bas. 1984. Belief and the Will. *The journal of philosophy*, **81**(5), 235–256.
- Varzi, Achille. 2007. 'Supervaluationism and its logics'. *Mind*, **116**(463), 633–676.
- Weatherson, Brian. 2003. 'From classical to constructive probability'. *Notre dame journal of formal logic*, **44**, 111–123.
- Williamson, J. 2010. *In defence of objective Bayesianism*. Oxford University Press.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, Timothy. 2006. 'Must do better'. In: Greenough, Patrick, & Lynch, Michael (eds), *Truth and realism*. Oxford: Oxford University Press.
- Wright, Crispin. 1992. *Truth and objectivity*. Cambridge, Mass.; London: Harvard University Press.
- Wright, Crispin. 2001. 'On being in a quandry'. *Mind*, **110**.
- Wright, Crispin. 2003. 'Vagueness: A fifth column approach'. *Pages 84–105 of: Beall, J. C. (ed), Liars and heaps: New essays on paradox*. Oxford: Oxford University Press.