THE REVIEW OF SYMBOLIC LOGIC
Volume 4, Number 1, March 2011

DEGREE SUPERVALUATIONAL LOGIC

J. ROBERT G. WILLIAMS
Department of Philosophy, University of Leeds

Abstract. Supervaluationism is often described as the most popular semantic treatment of indeterminacy. There's little consensus, however, about how to fill out the bare-bones idea to include a characterization of logical consequence. The paper explores one methodology for choosing between the logics: pick a logic that *norms belief* as classical consequence is standardly thought to do. The main focus of the paper considers a variant of standard supervaluational, on which we can characterize *degrees of determinacy*. It applies the methodology above to focus on *degree logic*. This is developed first in a basic, single-premise case; and then extended to the multipremise case, and to allow *degrees* of consequence. The metatheoretic properties of degree logic are set out. On the positive side, the logic is supraclassical—all classical valid sequents are degree logic valid. Strikingly, metarules such as cut and conjunction introduction fail.

§1. Introduction. Supervaluationism is often described as the most popular semantic treatment of indeterminacy. There's little consensus, however, about how to fill out the barebones idea to include a characterization of logical consequence. In a recent paper, Varzi (2007) writes:

> it is pretty clear that there is not just one supervaluational semantics out there—there are lots of such semantics; and although it is true that they all exploit the same insight, their relative differences are by no means immaterial . . . a lot depends on how a given supervaluational machinery is brought into play when it comes to explaining the logic of the language. (p. 463)

The 'supervaluational machinery' to be discussed here is the idea of a *supervaluational model*. Varzi highlights the fact that all sorts of properties of sequents that are candidates for the name 'validity' are available in this setting. This raises the question of how to choose among competing characterizations; and more fundamentally, what is at stake in the choice. One way to make progress is to make explicit the role of logical consequence within wider theory, and compare the impacts on wider theory that rival characterizations of consequence will have. We characterize what the *role* of logic is, and represent different candidate explications of logical consequence as reflecting different opinions about what properties play that role.

Section 2, below, introduces the basic machinery of supervaluational models, and connects issues over *choosing logic* to views on the *cognitive role of indeterminacy* (this connection is emphasized by Field (2003). The third section introduces a rival to Field's conception of the cognitive role of indeterminacy—a 'degreed' cognitive role. This motivates a particular kind of 'degree logic', whose properties are explored. The fourth

Received: July 04, 2010
Parts of the material were presented at the Leeds *CMM* group and at the Australasian Association of Philosophy annual conference in 2009.

© Association for Symbolic Logic, 2010

doi:10.1017/S1755020310000237

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

section explores extensions to this logic to multipremise inferences and inconsistency, and to *degrees of validity*. Appendices summarize some metatheoretic properties of the notions introduced, and prove key supraclassicality results for each.

§2. Supervaluational logic and norms for belief.

2.1. Supervaluational models. A basic supervaluational model $m = \langle D, S, I \rangle$ for a propositional language L contains of a set of 'delineations' D, and an interpretation function I that assigns T or F to each sentence in L relative to each delineation.[1] The model contains a set of delineations that are special—the sharpenings S. A sentence letter A is true-on-d, m iff I_m maps A to T relative to d. A is supertrue-on-m iff A is true-on-d, m for every sharpening $d \in S$. A sentence is supertrue simpliciter iff it is supertrue on the intended model. Bivalence fails for supertruth. If a sentence is true-on-d, m and false-on-d', m for $d, d' \in S_m$, then neither it nor its negation is supertrue-on-m. Whether truth itself should be identified with supertruth is not something we'll take a stance on at this stage.

Supervaluational models have found many kinds of application. Here are three:

2.1.1. Application 1: open future. Perhaps the delineations are big-bang-to-heat-death linear world histories. The sharpenings are the histories that are compatible with past-and-present facts, and the laws of nature. If the laws of nature are nondeterministic, there will be many historically possible evolutions of the universe. Correspondingly, there will be multiple sharpenings. What is supertrue is what is determined to be true by the laws and past facts. Those not so fixed—'future contingents'—are neither supertrue nor superfalse. This model has prima facie attraction for those who think that there are no 'future facts' to truthmake future contingent statements (e.g., think of some 'growing block' theories).[2]

2.1.2. Application 2: indeterminacy in the world. The delineations are complete histories as above, but we want to apply the story more widely. Take Kilimanjaro, and a small particle, Sparky, somewhere near its edge. Is Sparky part of Kilimanjaro? Maybe it's not determinately the case that it is, nor determinately the case that it isn't (where this indeterminacy isn't to be thought of as 'linguistic' but rather 'metaphysical'). Let the sharpenings be the histories that don't determinately misdescribe the situation—so one sharpening says that Sparky is part of Kilimanjaro (not determinately right, but not determinately wrong, either); and another says that Sparky is not part of Kilimanjaro (ditto). 'Sparky is part of Kilimanjaro' is neither supertrue nor superfalse.[3]

2.1.3. Application 3: semantic indecision. The delineations correspond to fully classical interpretations of our language. The sharpenings are those that correspond to interpretations that are not ruled out by the meaning-fixing facts—that don't conflict with conventional regularities of assent and dissent, and that are sufficiently 'eligible'.[4] If Harry is borderline bald, then perhaps there is no conventional regularity of assent to 'Harry is bald', nor to its negation—and so perhaps, there will be a sharpening on which this is

[1] The *admissible* interpretations are those that which respect the standard classical rules for the truth functional connectives. From now on I'll assume that all 'classical interpretations' are admissible.

[2] See for example Thomason (1970). This sort of model, with tweaks that leave unchanged the basic behavior, underlies the recent discussion of future contingents and relativism (Macfarlane, 2003).

[3] See for example Akiba (2000), Barnes & Williams (2010).

[4] Lewis (1984), Williams (2007).

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

true, and a sharpening where this is false. Such borderline sentences will then be neither supertrue nor superfalse.[5] The semantic indecision interpretation of supervaluationism is prominent in the literature on paradigmatically vague terms (like 'bald') but has also been invoked in several other areas: conditionals (Stalnaker, 1980); theory change (Field, 1973); and the problem of the many (Lewis, 1993).

In this paper, I will use these potential applications for illustrative purposes, but I won't be defending the view that supervaluational models are the right tool for the job. Our concern here is not primarily with issues of application, but with how supervaluational model theory interacts with logic and ideal (partial) belief.

2.2. Determinacy and degree determinacy. A supervaluational model will characterize some sentences as neither supertrue nor superfalse. On the *intended* or *designated* model, this gap is supposed to pick out those sentences that are *indeterminate* in the relevant sense. It is standard to allow the object language under consideration to express claims about what is or what is not indeterminate. We add to the propositional base language a sentential operator D we can gloss as 'it is determinate that'. The compositional clause needed is the following:

- $D\phi$ is true-at-i, m iff ϕ is true-at-s, m for all $s \in S_m$.

Notice that if A stands for any of 'Harry is bald', 'Sparky is part of Kilimanjaro' and 'The coin will land heads' in the above examples, $\neg DA \wedge \neg D\neg A$ will be supertrue—a statement which translates as 'it is indeterminate whether A'. Edgington (1997), Lewis (1970), and Kamp (1975) argue that we should add further detail to the above story.[6] In particular, rather than simply having an all-or-nothing distinction between those p which are determinately the case and those which are not, the suggestion is that we think in terms of *degrees of determinacy*. In the version in which I will consider it, the idea is that when p is true on *almost all* sharpenings, it is *more* determinate than a q which is true on *almost no* sharpenings.

We need to add a little more structure to the supervaluational models to capture this idea in the general case. Formally, we add to our supervaluational model a measure μ_m over the set of sharpenings.[7] What does this measure come from? There may be no uniform answer

[5] See for example Field (1974), Fine (1975), and Keefe (2000). Sometimes further formal constraints are imposed on the sharpenings by restricting the supervaluational models to those consistent with some underlying Kleene model. See the Fine article for this sort of restriction.

[6] The authors' proposals, and the aims they have for introducing them, vary.
None are explicitly committed to the cognitive role application that is our central focus here. Lewis and Kamp have primarily semantic concerns—giving a compositional treatment of comparative forms of adjectives and of modifiers such as 'rather' or 'very'. Edgington is focused on diagnosing and defusing the sorites paradox, and though explicitly committed to a probabilistic degree structure for determinacy, is wary of associating this with a supervaluational semantics.

[7] μ maps subsets of S into [0, 1]. By requiring it be a measure (rather than a 'finitely additive measure') we're assuming *countable additivity*. Strictly, a measure is defined over a measurable space—in this case it would be a pair of the set of sharpenings S and a σ-field of subsets of S. Subsets of S that are not in the σ-field would be 'nonmeasurable'. If S is infinite, we can't guarantee there won't be nonmeasurable sets. This leads to the question of how we handle the logic and semantics of sentence letters that are true exactly on a nonmeasurable set of sharpenings (cf. Smith, 2008, pp. 88–90). I'm in effect going to stipulate this worry out of existence somewhat, by restricting attention to *admissible* models of the language, where this is understood to require that the set of sharpenings making a given sentence letter true is required to be measurable (which means by finite additivity and complementarity, every Boolean combination of sentence letters

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

to this question; but local to each application, one might claim we can describe how an 'intended measure' could be determined. A natural example in the case of the 'open future' would be to identify the measure of X with the objective chance, at a particular time, of one among X being the actualized history. In the case of semantic indecision, we might think of the relative weight of each subset as fixed by how well the subset collectively 'scores' by the lights of the meaning-fixing facts. Sometimes the idea is presented as if what mattered were the *relative number of the total sharpenings that are within X*—in the finite case, just the size of X divided by the total number of sharpenings (whether such a 'pure counting' measure even makes sense when there are infinitely many sharpenings is contentious). The construction of such measures is likely to be philosophically controversial. But controversial or not, the logical issues developed below are unaffected by these differences in the 'metaphysics of the measure', so I set the issue aside and simply take an intended measure for granted.

Given this background, we can introduce degreed determinacy operators, allowing us to express claims such as 'the degree to which it is determinate that Harry is bald has increased over the years'. Here is the natural characterization:

- $D_{\geq k}\phi$ is true-at-i, m iff $\mu_m(X) \geq k$, where X is the set of sharpenings s such that ϕ is true-at-s, m.

We can then introduce a range of defined operators: for example the precise-degree determinacy operator $D_k p := D_{\geq k} p \wedge D_{\geq 1-k} \neg p$. It's easy to see that this will be true-on-s when p is true at exact measure k of the sharpenings. The operators, as presently characterized, have no nontrivial iterations. If it's true at one point in model that $D_{\geq k}\phi$, then it's true at all such points. So if $D_{\geq k}\phi$ itself determinate to any nonzero degree, it is determinate to degree 1. The operators are 'S5 like' in this respect. I would urge readers not to be overly concerned by this feature. For one thing, for certain applications it may be desirable. With respect to the supervaluational 'open future' described earlier, there's no clear motivation for significant iterated indeterminacy. In other applications, significant iterated indeterminacy may be desirable—iterations of nondegreed determinacy operators form a central part of the discussions of 'higher order vagueness', for example (though even here the relation between philosophical and formal questions is hotly contested[8]). While it would be possible to add to the complexity of our models in order to allow for nontrivial iterations, it seems methodologically preferable to study the simpler cases initially, with the aim of cautiously extending this core model if a particular application demands it.

Fans of the above semantic model can make sense of much of the talk about indeterminacy that 'fuzzy logicians' go in for.[9] For example, in application to vagueness, they can say that adjacent members of a sorites series for 'bald' differ in their degree of determinacy by only minute amounts—but that the minute differences add up over a sorites series, taking us from clear cases of baldness (perfectly determinate cases of baldness) to clear cases of nonbaldness (perfectly determinate cases of nonbaldness)—without any 'big

will be measurable). The proof of supraclassicality below appeals to this. A different salient option is to assign 'degrees of determinacy' not via the measure itself, but via the associated 'inner measure'. I will not pursue these issues further here, or explore its implications for the metatheory. With a single exception, in the discussion below we can restrict attention to the case where S is finite, S itself may be taken as the σ-field. A candidate (but not the only candidate) for μ is obtained by counting the number of sharpenings in a set.

[8] See the final chapter of Keefe (2000) for a defense of an S5 model of determinacy operators.

[9] See Machina (1976) and Smith (2008).

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

jumps' in status. Furthermore, they do this while avoiding what many find puzzling aspects of standard fuzzy logic 'degree theories'. On the supervaluational degree theory every classical tautology is true on each sharpening, and hence will be determinate to degree 1. Every classical inconsistency will be false on each sharpening, and hence determinate to degree 0. So while standard fuzzy logicians need to convince us that sometimes 'Harry is bald and not bald' can have a middling degree of determinacy, the current 'degree supervaluational' framework can say it has no nonzero degree of determinacy. Similarly for 'Harry is bald, and Jimmy is not' (where Jimmy is balder than Harry). As an optional extra, one can use the ideas just introduced to give precise sense to the idea of 'degrees of truth', characterized as follows:[10]

- ϕ is true to degree d (on m) iff $\mu_m(X) = d$, where X is the set of sharpenings s such that ϕ is true-at-s, m.

I'll help myself to this terminology when convenient, though strictly all talk of 'degree of truth' could be eliminated in favor of talk of degrees of determinacy.

2.3. Logic and cognitive role. There has been considerable debate among friends of supervaluational models about how to characterize logical consequence. There are many different perfectly legitimate properties of sequences of sentences one could look at (we'll see quite a few shortly). In certain situations, these diverge from each other. A suspicion naturally arises that the question 'which is the real logic' is without content—there are various generalizations of standard definitions, giving interestingly different results, and that's all we can say.

If logic was a theoretical idler—if it didn't *interact* with wider theory—then we could leave this here. But it is not an idler, I believe, and so it makes sense to press the matter further. The theoretical role for logic I FocUs on concerns *norming belief states*. For simplicity, let's start by thinking about 'all-or-nothing belief'. Someone who believes A and $\neg A$ has something pro tanto wrong with their state of belief. That's not to say that there is, all-in, some belief state that it would be better for them to adopt, especially if they have to deal with conflicting evidence. Nevertheless, inconsistency is a serious (not merely prima facie) flaw in a belief state. Here, then, is one logic–belief constraint: pro tanto, one's belief state shouldn't be inconsistent.[11]

I favor a representation of belief states via *partial belief* or *credence*, rather than all-or-nothing belief. But pro tanto logical norms still apply. Here is one, for example: one shouldn't have less credence in the conclusion of a valid sequent than in its premise. Note that this follows (and can be used as an axiom) in characterizing classical probability theory, and so the view that logic imposes such a pro tanto constraint is closely connected to the view that one's credences should satisfy the axioms of probability theory.

What sorts of characterizations of consequence might we look at? One thought is to rely on a standard gloss of logic as 'necessary truth preservation'.[12] However, we haven't

[10] There are marked differences between this and standard 'degree theories'—in particular, classical tautologies will have degree of truth 1, classical inconsistencies degree of truth 0, and there will be no function that settles the degrees of truth of logical compounds from the degrees of truth of their component parts. Edgington calls these notions 'verities' or 'degrees of closeness to truth'. Lewis and Kamp are happy with the terminology 'degrees of truth'.

[11] Maybe the constraint kicks in only when the inconsistency is or should be obvious—I leave this wrinkle to one side for now. See Field (2009) for discussion.

[12] See Williamson (1994, ch. 5) for an advocacy of such truth—logic connections.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

yet said what 'truth' in a supervaluational model is to be. A traditional supervaluationist answer (*the* supervaluationist answer, on one understanding of 'supervaluationist') is that truth is what we earlier called *supertruth*—truth at all sharpenings. So one characterization of consequence is 'supertruth preservation on every model', formally:[13]

- $\phi \models_g \psi$ iff there is no supervaluational model m on which ϕ is (super)true and ψ is not (super)true.

This 'global' notion of consequence is often contrasted with the following 'local' characterization:

- $\phi \models_l \psi$ iff There is no sharpening s and supervaluational model m on which ϕ is true at s, m and ψ is false at s, m.

If we restrict attention to the D-free fragment of our propositional language, these both coincide extensionally with a standard classical characterization of consequence for propositional languages. But in the D-including fragment the two characterizations come apart. A pair showing this is the following:[14]

- $p \models_g Dp$
- $p \not\models_l Dp$.

Setting aside issues of how to characterize truth, and the relationship between truth and consequence, what belief norms would global supervaluational consequence deliver? Suppose I'm fully confident that Harry is borderline bald. My credence that it's determinate that he's bald is vanishingly small. But as noted above, that it's determinate that he's bald is a g-consequence of his being bald. If g-consequence plays the belief state norming role, then my credence that Harry is bald shouldn't be greater than my credence that it's determinate that he's bald; and so I invest at most vanishingly small credence in his being bald. Mutatis mutandis for 'Harry is not bald'. Overall, the result is that my credence in 'Harry is bald', and its negation, should each be close to 0. In the limiting case, I should utterly reject each claim. Call this view of the cognitive role of determinacy 'rejectionism'. This is a striking consequence of taking $\models_g$ to norm belief. Notice, for example, that it conflicts with standard probabilistic models of credences, which insist that credences in p and in $\neg p$ should add up to 1. [15] The belief-norm role for logic connects the choice

[13] For now I focus on a single-premise, single-conclusion case since it is simpler to state, is adequate for the points to be made in the next few sections, and cuts down on the number of choice points in formulating a characterization of consequence.

[14] Many have noted that this result is called into question if we tweak the semantics to allow for a non-S5 logic for the D operator, by adding in an accessibility relation over the delineations. The consensus is that nonclassical behavior still surfaces, in particular in the pair $p \wedge \neg Dp \models_g$ and $p \wedge \neg Dp \not\models_l$. I discuss this latter claim in Williams (2008). For simplicity, I will stick with S5 models of determinacy throughout (for a defense of this as the right logic for D, see the final chapters of Keefe, 2000). This also famously forms the basis for the notorious failure of classical metarules for $\models_g$ (Williamson, 1994, ch. 5). Even though $p \models_g Dp$, we have the following: $\not\models_g p \supset Dp$. For the latter to be g-logically true, it'd have to be true on every supervaluational model. But a model with two sharpenings, on one of which p, and on the other of which $\neg p$, is one where Dp is false on both sharpenings—and hence the material conditional is false on the first. This is a supervaluational model where the conditional is neither true nor false.

[15] A formal model of credences appropriate to rejectionism (suggested by Field) can be found in the literature on Dempster—Shafer belief functions. See Halpern (1995). In (Williams, manuscript) I discuss these generalizations in a very general setting, and adapt the 'accuracy' arguments of

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

of logic with one's views on the *cognitive role of indeterminacy*—the rational constraints imposed by a judgment that p is indeterminate on one's wider beliefs state (including, in particular, one's degree of belief in p itself). I take it those sympathetic to the rejectionist cognitive role just characterized will favor the global consequence characterization at least over its local rival—for the local rival *fails to capture* rationality requirements that for the rejectionist *should be imposed* by a logic of determinacy. On the other hand, if we think determinacy should not have a rejectionist cognitive role, then we must reject $\models_g$ as a logic that norms belief states. But what alternative might we consider? My main focus in the remainder of this paper is on a logic that becomes available once we have measures in our supervaluational models:[16]

degree-preservation $\phi \models_{deg} \psi$ iff there's no m such that $\mu_m(\phi) > \mu_m(\psi)$.

Much of the remainder of this essay will be spent exploring the formal behavior of this logic and extensions thereof. But in order for this investigation to have more-than-formal interest, we need to say something about why one might select *this* characterization from among the multitude of possibilities. I will argue, in the section that follows, that there is a conception of the cognitive role of indeterminacy—a rival to rejectionism—on which the degree-logic norms belief.

§3. Degree supervaluationism. Rejectionism, as I see it, is one answer to a mandatory question for the theorist of indeterminacy. Suppose one judges that p is indeterminate. How should that impact your first-order opinion as to whether p or $\neg p$? The rejectionist has an answer: you should reject each of them utterly. But a rival thought is that faced with a smack-bang borderline case of baldness (like Harry) your credences should be 50/50.

Joyce (1998) to tie together model theory, logic, and probability in the generalized setting. More on this below.

16 This is one among a host of options. Even without appeal to measures, there are all sorts of options—see Varzi (2007) for a systematic discussion. One option is guaranteed preservation of nonsuperfalsity—which generates a 'subvaluational' logic. Another requires that both supertruth and nonsuperfalsity be preserved. There are nonrejectionist cognitive roles that can be associated with these logics, but I'm not going to explore these here. When we allow ourselves appeal to measures, the number of options that we may consider explodes. To bring some order, we can look at the characterizations that have proved interesting in a literature that's analogous to the present discussion: 'probability logic' (Adams, 1996). As well as the 'degree preservation' option in the text, this literature gives a prominent role to logics that in our setting become:

1-preservation $\phi \models_1 \psi$ iff there's no m such that $\mu_m(\phi) = 1$, and $\mu_m(\psi) < 1$.

Bound-preservation $\phi \models_b \psi$ iff for any $\delta > 0$, there is an $\epsilon > 0$, such that for any m, if $\mu_m(\phi) > 1 - \epsilon$ then $\mu_m(\phi) > 1 - \delta$.

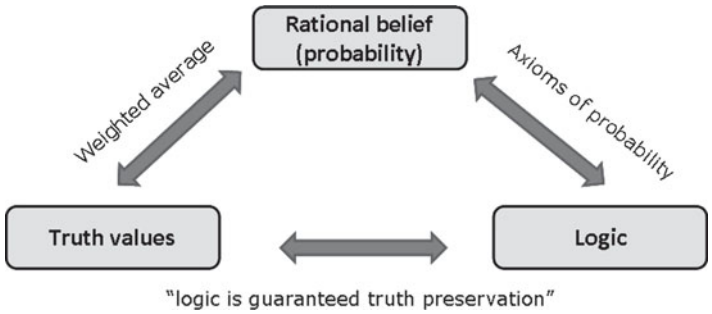
Adams' results show that in a determinacy operator—free fragment of language, these characterizations (and the degree logic) extensionally coincide. However, one can quickly check that they diverge when determinacy operators are added (a situation familiar from the global/local consequence debate). For example, $p \models_1 D_1 p$, but $p \not\models_b D_1 p$ and $p \not\models_{deg} D_1 p$. Also $p \models_b D_{\geq 0.999} p$, but $p \not\models_{deg} D_1 p$. These results can be used to show that bound and 1-preservation violate the classical metarule of conditional proof, in the same way that global consequence does. More significantly for our purposes, if we treat them as norming credence in the standard way, they enforce rejectionism just as global logic does. So I set them aside for present purposes.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

Clearly that's incompatible with rejectionism. However, denying rejectionism doesn't obviate the need to give some answer or other to the cognitive role question identified above—one minimally needs to explicate the obvious fact that answering the question 'p?' with the answer 'p is indeterminate' is a *relevant* response. What sort of answer should a friend of supervaluational models opposed to rejectionism give?

Here, the measures of sharpenings come into their own. If Harry is a smack-bang borderline case of baldness, one can expect that the proportion of sharpenings on which he counts as bald be 0.5 (that's a good way to understand what it is for someone to be a 'smack-bang' borderline case). Now consider what happens as Harry ages and loses more hair. Many feel one should smoothly increase your credence that he is bald. This is nicely matched by the degree of determinacy; the measure of those sharpenings on which Harry counts as bald can be expected to increase with hair loss. Or consider the case of future contingents in the supervaluational understanding of the 'open future' described earlier. Fair and biased coins are paradigmatic application points for credence. You should be 0.5 that a fair coin will land heads; 0.999 that a massively biased coin will land heads. On a natural articulation of the view, this matches up with the degrees of determinacy in each case. So a rival to rejectionism puts proportions of sharpenings into a central role. Credence (from a gods' eye point of view) should match the degrees of determinacy.[17]

3.1. Truth values, generalized probability, and logic. Thus far, I've characterized the degree cognitive role for indeterminacy in terms of *ideal aims* for belief. I've not said anything directly about the cognitive role of indeterminacy *judgments*. To appreciate the structure of the debate, let us consider the interrelations between formal representations of truth values, rational belief, and logic, as captured in this diagram:



Here, the nodes are all familiar. Truth value distributions are the standard bivalent distributions—representable as the assignment of 1 to true sentences, and 0 to false sentences; rational belief states are assumed to be described by probability functions; and classical logic fills out the picture. But the nodes are interrelated. Most familiarly, logic can be model theoretically characterized in terms of '1-preservation on every classical model'—connecting up distributions of truth values with logical consequence. This is the bottommost connection. Every classical probability can be represented as a *weighted average* of classical truth value distributions.[18] And completing the picture, it is possible

[17] This is a familiar thought in the literature on nonsupervaluational degrees of truth—see Smith (2010) and Macfarlane (2010).

[18] More precisely, concentrating on the finite case, we can suppose that the doxastically possible worlds are w_i, for $i \leq n$. σ is a probability assignment over a propositional language iff for every

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

to give 'local' characterizations of probability functions by a set of axioms that appeal, in one way or another, to logical relations: tautologies must be assigned probability 1, inconsistencies probability 0; when $A \models B$, $P(A) \leq P(B)$; and when A and B are inconsistent, $P(A \vee B) = P(A) + P(B)$.

This is an elegant and unified package. Notice how it embeds the normative significance for logic. For if we accept the characterization of rational belief states as probabilities, then because all probabilities satisfy certain logical constraints (as captured by the axioms) belief states that violate those axioms will be (at least pro tanto) objectionable.

In a nonclassical setting like that of the supervaluationist, by characterizing the 'ideal aim' of credence we in effect have a new way of filling in bottom left node. If we take the 'truth value' of a sentence to coincide with the credence one should ideally have, then granted the degreed conception, truth values will be degrees of truth, and can take values strictly between 1 and 0. The question then is—how do we fill in the rest of the picture? And in particular—what goes in the 'logic' slot?

Here, in short, is my favored answer (justification to follow!) Rational belief states should be characterized as *generalized* probabilities—where these are related to truth value assignments exactly as classical probabilities do to classical ones. That is, a generalized probability is exactly a weighted average of degree truth value assignments.[19] That fills in the top node and leftmost connection. We then look for what 'consequence relation' stands to these generalized probabilities as classical consequence stands to classical probabilities. And, I claim, this role is played by degree logic defined earlier. The characterization of degree logic as a 'guarantee of no drop in truth value' fills in the final connection.

What justification might we offer for this package? The most crucial, and controversial, step, is the description of rational belief states in terms of generalized probabilities. We should immediately note that generalized probabilities are intuitive plausible candidates for rational credences nonclassical setting. The basic idea might be put as follows: one first divides one's total confidence between various completely specified possibilities. One then looks, for each possibility, what one's credence in p would *ideally* be *supposing that possibility were actual*. And one's overall view of p is then got by mixing together these hypothetical 'ideal' credences, in proportion to the confidence invested in the relevant possibility. This produces classical probabilities if truth value distributions are classical, distinctly nonclassical results given a rejectionist conception of truth value distributions—and gives the recipe described above for the supervaluational degrees of truth we are considering.[20] We cannot expect that a defense of the rational requirement that beliefs be generalized probabilities will be any less controversial than the corresponding claim in the classical case; and I haven't got space to fully explore this now. However, it is worth noting that famous arguments for classical probabilism do appear to generalize. Paris (2001) proves a general form of the classic Dutch Book argument, showing that an agent will be 'Dutch Bookable' iff their credences fail to be a generalized probability in the sense just

---

p, $\sigma(p) = \sum_{i \leq n} \lambda_i w_i(p)$ where the 'weights' λ_i, are such that $\sum_{i \leq n} \lambda_i = 1$, and $w_i(p)$ is 1 iff p is true at w_i, and is 0 otherwise.

[19] That is we let $w_i(p)$ now pick out the degree of truth of p at w_i, and otherwise let the characterization stand as before: $\sigma(p) = \sum_{i \leq n} \lambda_i w_i(p)$ where the 'weights' λ_i, are such that $\sum_{i \leq n} \lambda_i = 1$.

[20] In Williams (manuscript) I examine this notion of generalized probability for a range of nonclassical views on truth values, including examining the rejectionist take in detail. Generalized probabilities in that setting turn out to be Dempster Shafer functions. Compare Field (2000).

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

given. Williams (manuscript) adapts Joyce's (1998, 2009) *accuracy domination* argument to argue that agents' belief states that are not generalized probabilities will be 'accuracy dominated'—that is, there will be a specific generalized probability that is guaranteed (no matter what the world is like) to be closer to the truth values than the belief state one started with.[21]

Though detailed discussion must be reserved for another occasion, the above motivates an assumption I will make in the rest of the paper—that given degree supervaluationism, any belief state that is not a generalized probability is (pro tanto) irrational. Holding this fixed, we turn to the justification of the final element of the jigsaw. In the classical case, classical probabilities ρ may be characterized (with some redundancy) by the following set of schema, where $\models$ is read as classical consequence:[22]

P0 If $A \models$ then $\rho(A) = 0$
P1 If $\models A$ then $\rho(A) = 1$
P2 If $A \models B$ then $\rho(A) \leq p(B)$
P3 $\rho(A) + \rho(B) = \rho(A \vee B) + \rho(A \wedge B)$.

It turns out that *exactly the same set of principles* hold good for our generalized probabilities, if $\models$ is read via the degree logic.[23]

This is far from automatic. Global consequence, as one would expect, flunks out. Requiring credence not to drop over globally valid sequents (P2, above) would rule out perfectly decent generalized probabilities (e.g., those that assign high credence to $D_{0.5}p$, intermediate credence to p, and low credence to Dp). Logics *weaker* than degree-logic (such as local consequence[24]) won't suffer from this 'overgeneration' problem. But if they fail to declare valid degree-valid sequents, then they will *undergenerate* normative constraints on belief, as we've seen above that when $p \models_{deg} q$, every generalized probability will have

---

21 Both Paris's result and mine show that *quite* generally, belief states that are not mixtures of truth value assignments are subject to these results—the degree supervaluational version of particular interest here being a special case.

22 These axioms are taken from the discussion of nonclassical probability in Weatherson (2003), and are discussed at length in Williams (manuscript).

23 (P2) is perhaps the most crucial result for our purposes. Here is a proof. If B is a generalized probability then there are λ_i such that for each p, $B(x) = \sum_i \lambda_i w_i(x)$. In particular, $B(p) = \sum_i \lambda_i w_i(p)$ and $B(q) = \sum_i \lambda_i w_i(q)$. Ex hypothesi, $p \models_{deg} q$, so by construction for each w_i, $w_i(p) \leq w_i(q)$. Substituting this into the equations above we have $B(p) = \sum_i \lambda_i w_i(p) \leq \sum_i \lambda_i w_i(q) = B(q)$. Hence when p follows from q, and $B(p) > B(q)$, then B is not a generalized probability.

24 To see local consequence is weaker, note that if on every sharpening in a model, if A is true than B is, then a fortiori the set of sharpenings on which A is true is a subset of the set on which B is true, and hence the degree of truth of A, on that model, is no greater than that of B. We can use this to show that every generalized probability will satisfy the axioms of 'classical' probability (so long as we regard local consequence as the appropriate analogue of classical probability for a language containing the D operators, which is a plausible assumption). (One might have expected this in the present setting—after all, a generalized probability is a mixture of classical probabilities—and mixtures of classical probabilities are classical probabilities. But it doesn't hold in other nonclassical settings, where mixtures of nonclassical truth value distributions can lead to highly nonclassical conceptions of probability.) As noted in the text, the converse need not hold—not everything that meets the axioms of 'classical probability' so construed need be a generalized probability.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

$B(p) \leq B(q)$.[25] In the respects mentioned, degree logic stands to rational beliefs in our nonclassical setting as classical logic stands to rational belief in the probabilistic setting.[26]

3.2. Some metatheoretic results for degree logic. I have been arguing that someone with a 'degreed' conception of the cognitive role of determinacy judgments should think that the logic that norms belief-states is the degree-logic characterized above. Global supervaluational logic stands in a notoriously subtle relationship to classical logic—validating classical tautologies, but producing counterexamples to classical metarules. It's worth taking a snapshot of how degree logic behaves. At the moment, we are restricting attention to single-premise, single-conclusion sequents (generalizations and proofs are given in the appendices).

The central positive result is that degree logic (like global supervaluational logic) is **supraclassical**. That is, whenever $\phi \models \psi$ is classically valid (where all formulae are D-free), we have $\phi \models_{deg} \psi$. This follows from the stronger (multipremise) proof given in an appendix. Otherwise, the consequence relation has many properties one would expect—it is reflexive, transitive. It is worth noting that single-premise **reductio** holds, which is not the case for global consequence[27] (if $p \models_{deg}$, then we know that on any supervaluational model, the measure of sharpenings on which p is true is 0, which tells us that the measure of sharpenings on which $\neg p$ is true is 1).

The most dramatic result at this stage, without precedent in global consequence, is the failure of **conjunction introduction**. To see this, consider the following pair (recalling that $D_{0.5}p$ is true at a model iff exactly half the sharpenings make p true):

$$p \wedge D_{0.5}p \models_{deg} \neg p$$

$$p \wedge D_{0.5}p \models_{deg} p.$$

Consider an arbitrary supervaluational model. Either p is true on exactly half the sharpenings of that model, or it isn't. If the former, then the degree of truth of $D_{0.5}p$ is 1, the degree of truth of p is 0.5, and hence the degree of truth of the premise is 0.5. Clearly this is no greater than the degree of truth of p, nor of $\neg p$, which are each equal to 0.5 on this scenario.

If the latter scenario is realized, and p is not true on exactly half the sharpenings, then the degree of truth of the premise is 0 since the degree of truth of one conjunct ($D_{0.5}p$) is 0. In this case trivially it is no greater than the degree of truth of p, nor of $\neg p$.

So each of the above sequents is valid in degree logic. But conjunction introduction (in metarule form) would take us from $A \models B$ and $A \models C$ to $A \models B \wedge C$. Applied to this

[25] Here's an illustration of undergeneration in the local setting. Consider the sequent $D_1 p \models p$. This sequent holds on the degree-preservation logic. However, it isn't valid on the local characterization of consequence. It might be that $D_1 p$ is true on s, m (i.e., p is true on measure 1 of the sharpenings of m), while p is false at s, m (measure 1 of sharpenings needn't be *all* the sharpenings—at least if there are infinitely many of them).

[26] This is no accident—in Williams (manuscript) I show that there's a *general* recipe for characterizing a logic given a specification of 'truth values' or 'ideal aims for credence' that will play this role in axiomatizing the relevant notion of generalized probabilities. Degree logic is what the recipe delivers in this particular case.

[27] Williamson's example is that $p \wedge \neg Dp \models_g$, but $\not\models_g \neg(p \wedge \neg Dp)$—the latter would mean that $p \supset Dp$ was a tautology, which would clearly be disastrous.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

particular case, this would be:

$$p \wedge D_{0.5}p \models_{deg} p \wedge \neg p.$$

But this is *not* valid in degree logic. As we've seen, it's possible for the premise to be 0.5 true. But the conclusion is a classical contradiction, and will be 0 on every supervaluational model. So we have a drop in truth value. Conjunction introduction fails.

The formal features of degree logic are an interesting variation on the themes familiar from investigations of global logic. Supraclassicality shows that it behaves classically in the language free of D-operators. When D-operators are added, classical metarules fail—though we've seen above that some metarules that fail on global logic are preserved here (e.g., single-premise reductio); but others that global logic preserves fail here (e.g., conjunction introduction).

§4. Extending degree supervaluational logic. So far, we've considered only the single-premise case of degree logic. This may seem overly restrictive. Surely we want to be able to evaluate whether theories are consistent or not, and discuss whether rules like modus ponens ($A, A \supset B \models B$) or a sequent form of conjunction introduction ($A, B \models A \wedge B$) hold—the arguments we deal with are rarely restricted to a single premise![28]

But there are also motivations for considering extensions to the single-premise case that spring directly from our stated ambition for a characterization of logic in norming belief. Though we have shown that every generalized probability satisfies P0–P3, we have not shown that any function satisfying P0–P3 is a generalized probability.[29] So for all we've said, the normative constraints imposed by single-premise degree logic (in conjunction with other standard constraints) may undergenerate rational constraints.[30] This seems to me a possibility to be taken very seriously. For example, it is easily checked that on all generalized probabilities, if the probability of $D_{0.5}p$ is 1, then the probability of p must be 0.5. No rational belief state, therefore, should assign both of them, or their conjunction $p \wedge D_{0.5}p$, a probability of 1. But I do not see anything to rule this out in P0-P3. In the following two sections, I will consider two extensions to the degree-logic developed above—incorporating first *multiple premises* and second *degrees of validity*. I will formulate corresponding normative constraints, relating them to the example just given. This exercise is exploratory, and while I do not have a completeness proof to offer, I hope to have provided the resources that will be needed.

---

[28] For full generality, we might also wish to generalize to the case of multiple conclusions—see Varzi (2007), Restall (2005).

[29] Contrast standard 'completeness' results for classical probabilities. Similar results for a range of nonclassical cases for 'generalized probabilities' in our sense are presented in Paris (2001).

[30] It's worth noting that the more usual formulation of P3, which requires that $P(A) + P(B) = P(A \vee B)$ in situations where A and B are *inconsistent*, would indeed require multipremise validity. The inclusion–exclusion form used above finesses this. It's worth noting that axioms like $P3$ are 'impure', to the extent that they presuppose specific behavior of conjunctions and disjunctions (contrast with the other axioms, which only appeal to the logic to constrain probabilities). It would be elegant if we could formulate a 'pure' form of probability axioms that are entirely framed in terms of consequence. There is a way to do this (a common generalization of $P2$ and $P3$) but it requires not only multipremise consequence but a rather nonstandard version of multiconclusion consequence, and would take us too far afield.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

4.1. Multiple premises. Here's a natural way of understanding multipremise consequence (for finite premise sets, which we'll restrict ourselves to here for simplicity). On this proposal, we reduce multipremise consequence to the single-premise case, via conjunction:

conjunctive degree preservation

$\Gamma \models_{conjdeg} \psi$ iff $\bigwedge_{\gamma \in \Gamma} \gamma \models_{deg} \psi$.

Accordingly, a set Γ will be inconsistent, if and only if the conjunction of its members entail everything in the single-premise degree logic. This is a simple, easy to understand proposal. It certainly gives us *something to say* in analyzing arguments. I'll urge, however that an alternative (and extensionally divergent) characterization can do useful work, particularly when it comes to articulating constraints on belief.

Let us consider, therefore, the normative role of inconsistency in the classical case. One thing to note is that if a set of sentences Γ is inconsistent, this is not rationally incompatible with having high credence in each member of Γ. Γ could consist, for example, in claims about a fair lottery. For each n, Γ includes the claim that ticket n will lose; and also contains the claim that one of the tickets between 1 and 1000 will be a winner. One can rationally be highly confident in each member of this classically inconsistent set. However, one cannot rationally be *completely* confident in each member of Γ, however—one must be at least a little uncertain about enough of the claims. What is required is that *collectively* enough uncertainty is distributed. This feature is captured by the following connection between classical inconsistency and probability:

If $\Gamma \models$, then $\sum_{\gamma \in \Gamma}(1 - c(\gamma)) \geq 1$.

There is a nice principle that generalizes both the above and our earlier single-premise constraint (that credence shouldn't drop over a valid argument):

If $\Gamma \models \phi$, then $\sum_{\gamma \in \Gamma}(1 - c(\gamma)) \geq 1 - c(\phi)$.

In words—any unconfidence you have in the conclusion must be distributed through the premises, if the argument is valid.

We noted earlier that one's total confidence in $p \wedge D_{0.5}p$ cannot be greater than 0.5 (on a generalized probability). A related phenomenon concerns the following pair:

$$p,\ p \wedge D_{0.5}p.$$

It turns out that the sum uncertainty in these two sentences on any generalized probability must be at least 1.[31] So this pair is 'inconsistency-like' when it comes to rational credence. They're certainly not inconsistent by the conjunction criterion introduced above, however—it's easily checked that $p \wedge p \wedge D_{0.5}p \not\models_{deg}$. This motivates a different way of characterizing finite multipremise consequence, drawing on a characterization used in the literature on probabilistic logic.[32] First, define the degree of *un*determinacy (lack of determinacy) of p in a model m, as the measure of the sharpenings on which p is false in that model. Let's write this as $u_m(p)$. The 'summing' conception of multipremise validity

[31] A generalized probability ρ satisfies $\rho(x) = \sum_i \lambda_i w_i(x)$. On each w_i, the truth value of the conjunction $p \wedge D_{0.5}p$ is either exactly 0.5, or it is 0. Let's suppose worlds of the former sort have weight k. Then the credence in the conjunction is $(1/2)k$, and hence its uncertainty is $1 - (1/2)k$. Credence in p, on the other hand, is bounded above by $(1-k)+(1/2)k$, and hence the uncertainty in p is no less than $(1/2)k$. The sum uncertainty is therefore no less than $1-(1/2)k+(1/2)k = 1$.

[32] Adams (1996)—where this is used as the definition of 'probabilistic validity'.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

is that the consequence relation obtains when there is a matter of a guarantee that the *total* undeterminacy doesn't increase:[33]

summing degree preservation

$\Gamma \models_{sumdeg} \psi$ iff on every model m, $\sum_{\gamma \in \Gamma} u_m(\gamma) \geq u_m(\psi)$.

The corresponding notion of inconsistency is that the sum undeterminacy of the premises is at least 1.

Turn back to the pair given above: $\{p, p \wedge D_{0.5}p\}$. This pair is inconsistent on the summing characterization.[34] It's worth noting two points about this pair. First, the summing characterization is a proper extension of the conjunctive characterization—every sequent that is conjunctively valid is also summing valid; though as we've seen, not vice versa.[35] Second, summing validity violates the structural rule **Cut**, which tells us we can move from $\Gamma, \phi \models \psi$ and $\Gamma \models \phi$ to $\Gamma \models \psi$. We've just seen that $p, p \wedge D_{0.5}p \models_{sumdeg}$, and (as we saw earlier) $p \wedge D_{0.5}p \models_{sumdeg} p$. But $p \wedge D_{0.5}p \not\models_{deg}$. Some metalogical properties of summing degree preservation logic are given in Appendix A.

We earlier stated a constraint on rational credence that held in the classical case. The analogue is satisfied by every generalized probability, using the summing characterization:[36]

If $\Gamma \models_{sumdeg} \phi$, then $\sum_{\gamma \in \Gamma}(1 - c(\gamma)) \geq 1 - c(\phi)$.

If we are allowed to appeal to the above connection between multipremise validity and rational belief as an extra normative constraint, might we make progress? If we allow *repeated* premises we can indeed make such progress. For example, not only is $p, p \wedge D_{0.5}p$ inconsistent by the above criterion, but also so is the pair $p \wedge D_{0.5}p, p \wedge D_{0.5}p$.[37] And this fact will then tell us that $2.(1 - c(p \wedge D_{0.5}p)) \geq 1$, which simplifies to $c(p \wedge D_{0.5}p) \leq 1/2$. So the multipremise validity *characterized via summing validity* shows promise in capturing the normative constraints of logic on rational belief.

4.2. Degrees of validity. The second extension to degree logic I want to outline is prefigured in a rather different setting by Machina (1976). This is a rather more direct and general take on formulae such as $p \wedge D_{0.5}p$. Multipremise validity, on the summing conception, *indirectly* imposed constraints, by finding associated formulae that were 'fully

---

[33] It is a famous result in probability logic that this kind of summing characterization coincides with classical multipremise validity *for a language involving no probability or determinacy operators*. See Adams (1996) for the proofs in the context of probability logic, and the proof of supraclassicality in the appendix below for a formulation in the present context.

[34] In those models where p is 0.5 true, $p \wedge D_{0.5}p$ is true to degree 0.5, so $u(p) + u(p \wedge D_{0.5}p) = 1$. In all other models, $u(p \wedge D_{0.5}p)$ alone is 1, so $u(p) + u(p \wedge D_{0.5}p) \geq 1$. So $\{p, p \wedge D_{0.5}p\}$ is inconsistent by the criterion just given. (Thanks to Jason Turner for pointing this out to me).

[35] Proof: if $\alpha_1, \ldots \alpha_n \models_{conjdeg} \phi$, we have by construction $\alpha_1 \wedge \ldots \wedge \alpha_n \models_{deg} \phi$, and hence $\alpha_1 \wedge \ldots \wedge \alpha_n \models_{sumdeg} \phi$. But $\alpha_1, \ldots \alpha_n \models_{sumdeg} \alpha_1 \wedge \ldots \wedge \alpha_n$ and hence by transitivity of sumdeg (see later) we have $\alpha_1, \ldots \alpha_n \models_{sumdeg} \phi$

[36] Start with the LHS. $\sum_{\gamma \in \Gamma}(1 - c(\gamma))$ reduces to $\sum_{\gamma \in \Gamma}(1 - \sum_i \lambda_i w_i(\gamma)) = \sum_{\gamma \in \Gamma} \sum_i (\lambda_i - \lambda_i w_i(\gamma)) = \sum_{\gamma \in \Gamma} \sum_i \lambda_i (1 - w_i(\gamma))$. And this in turn can be written as $\sum_{\gamma \in \Gamma} \sum_i \lambda_i (u_i(\gamma)) = \sum_i \lambda_i \sum_{\gamma \in \Gamma}(u_i(\gamma))$, where $u_i(\gamma)$ is the degree of undeterminacy of γ at world i. Similarly $1 - c(\phi) = \sum_i \lambda_i u_i(\phi)$. But we know from the characterization of summing indeterminacy that for each i, $\sum_{\gamma \in \Gamma} u_i(\gamma) \geq u_i(\phi)$, and hence each element of the sum $\sum_i \lambda_i [\sum_{\gamma \in \Gamma}(u_i(\gamma))]$ is $\geq$ than the corresponding member of the sum $\sum_i \lambda_i [u_i(\phi)]$. This shows that the LHS $\geq$ RHS.

[37] Though $p \wedge D_{0.5}p$ alone is not, of course—summing validity is not *idempotent*—repeated premises cannot be eliminated while preserving validity.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

inconsistent' by its criterion. But one might think that the *obvious* thing to say of this case is that it itself is *partially* inconsistent. A formula is fully inconsistent, in the (single-premise) degree logic if it is guaranteed to have a degree of truth no greater than zero. A formula is *partially* inconsistent if it is guaranteed to have a degree of truth no greater than k, for some $k : 0 < k < 1$.

The suggestion is, therefore, that one can view degree validity as the limiting case of a whole family of degreed notions of validity. Our earlier degree logic is then the 0-indexed version, which we can write as follows:

- $\phi \models^{0}_{deg} \psi$ iff there's no m $\mu_m(\phi) > \mu_m(\psi)$.

That is: validity is guaranteed preservation of degree of determinacy. The generalization is the following:

- $\phi \models^{k}_{deg} \psi$ iff there's no m $\mu_m(\phi) > \mu_m(\psi) + k$.

That is: k-validity obtains where we're guaranteed that any fall in degree of determinacy over the sequent is *no greater than k*.

The natural generalization of the normative constraint of (0-indexed) degree logic is this: pro tanto, one's credences shouldn't drop over an argument more than k where degree of determinacy is guaranteed to drop by no more than k. The degree of validity of an argument imposes rationality constraints on belief states, by requiring differences in credence to be no more than the degree of validity of the corresponding argument.

Notice that though $p \wedge D_{0.5}p \models^{0}_{deg}$ doesn't hold, $p \wedge D_{0.5}p \models^{0.5}_{deg}$ does. For we know that the conjunction is either determinate to degree 0, or determinate to degree 0.5. And this means that no other claim can possibly be less determinate than the premise by more than 0.5—just as required. The normative constraint then correctly captures what's wrong with fully believing $p \wedge D_{0.5}p$. For presumably, one fully disbelieves absurdities—and yet those absurdities follow from this conjunction to degree 0.5. By the logic—credence link, it's pro tanto wrong for the difference to be more than 0.5 between these your credence in $p \wedge D_{0.5}p$ and your credence in the absurdity. Having credence 0 in absurdities, and more than 0.5 credence in the conjunction, violates this constraint.[38]

§5. Conclusion. We started from an initial characterization of *supervaluational models*. Appeal to structures of this kind are widespread in philosophy—in the philosophy of time, in paradigmatic vague predicates, in recent work on metaphysical indeterminacy, and elsewhere.

But appeal to such models doesn't settle the question of what the best way to characterize the logic for the relevant languages is. I've related this to the question of what the *cognitive* or *doxastic* role of indeterminacy is. I've been investigating a conception on which one's strength of belief in p should (ideally) match the proportion of sharpenings on which p is true. The question arises: what sort of characterization of logic should we give, if we want it to play the normative role just specified? Many familiar candidates overgenerate—claiming rational constraints that simply are not there on the conception of indeterminacy in question. I urge degree logic as the best candidate, and pointed out some of its (rather interesting) metatheory. Our whole discussion has proceeded on the assumption that the

[38] One can combine degrees of validity with the multipremise validity in the obvious fashion—I won't spell out the details here.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

logic we're interested in is one that norms belief states. I have no argument, for example, that the summing degree logic is what we need to think about to account for appropriate direct and indirect reasoning.[39] One of the nice things about the traditional classical package is that, arguably, the indirect-reasoning applications of logic and the belief-state norming applications are discharged in an elegantly unified way. It's not so clear that in nonclassical settings the same holds true. This shows again the value of making explicit what sort of question we're asking, when we ask what the correct logic is for nonclassical treatments of indeterminacy.[40]

In the final sections, I consider ways of extending the core single-premise version of degree logic. This is motivated if—as seems likely—our original characterization itself undergenerates rational constraints. I've illustrated how generalization to the multipremise case, or introduction of degrees of validity, can capture normative constraints on credence that were previously problematic.

§6. Acknowledgment. I'd like to thank all with whom I've discussed the material, including Elizabeth Barnes, Ross Cameron, Daniel Elstein, Carrie Jenkins, Daniel Nolan, Jason Turner, and Nick JJ Smith.

Appendix

A Summary of the properties of sumdeg validity. In an earlier section, I ran through some metatheoretic properties of the single-premise degree logic. For completeness, this appendix extends these results to the multipremise and degree validity settings. A key result is proved in Appendix B. Our first set of results concern sum-degree validity:

1. It is **Supraclassical** for (finite) multipremise sequents. That is, whenever $\Gamma \models \phi$ is classically valid (where Γ is finite and all formulaes are D-free), we have $\Gamma \models_{sumdeg} \phi$. The proof of this is given in Appendix B.
2. It is **Reflexive**. The consequence relation has the obvious property that if $\phi \in \Gamma$, then $\Gamma \models \phi$. For the sum undeterminacy of the premises in this case obviously

[39] It's interesting that Adams (1996), in discussing probability logic, deploys a Hilbert-style proof theory that does not allow for indirect proofs. He makes the nice point that the philosophical interpretation of indirect proofs where some sentences have 'probabilities' alone rather than truth values, is problematic. A similar point could be made here—what sort of status is it that we attribute to a claim when we suppose it to be true?

[40] A final note on this point. In Williams (2008), it is noted that the failure of classical metarules for global consequence doesn't preclude the obtaining of *weakened versions*. (A similar point is made by McGee and McLaughlin, and in recent work by John Macfarlane.) For example, though conditional proof fails, so long as $\Gamma, A \models_l B$, we know that $\Gamma \models_g A \supset B$. This is no accident. For since conditional proof works for local logic (as local logic essentially just is classical logical consequence), the local version of the second sequent must hold. Then, by supraclassicality of global consequence, we get that the global version of the second sequent holds. The reasoning generalizes to every classical metarule: so long as the inputs to the metarule are all locally valid, the output will be (locally, and so) globally valid. The point carries over to our setting. Though metarules may fail in strange ways, the weakened variants where the input sequents are locally valid will guarantee the (local and thus) sum-degree validity of the outputs. While the metatheory may be structurally messy, there's little danger of the *applications* of the metarule that we're familiar with being problematic. So it's not as if *instances of reasoning* that we previously would have thought were in good standing are suddenly declared unreliable, so long as they fit the pattern of the weakened versions of the metarules.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

has to match or exceed the undeterminacy of the summands—which in this case includes ϕ.

3. It is **Monotonic**. Any valid sequent remains valid when you add extra premises. Adding a premise to an inference can only increase the sum total undeterminacy of the premises overall (in any model)—since validity requires the sum total undeterminacy be bigger than that of the conclusion, adding premises won't undermine a valid argument.
4. **Reductio** holds. Proof of reductio: Suppose $\Gamma, \phi \models_{sumdeg}$. That requires for arbitrary m, $\sum_{\gamma \in \Gamma} u_m(\gamma) + u_m(\phi) \geq 1$. This is equivalent to $\sum_{\gamma \in \Gamma} u_m(\gamma) \geq 1 - u_m(\phi)$. But $1 - u_m(\phi) = u_m(\neg\phi)$, and hence we have $\sum_{\gamma \in \Gamma} u_m(\gamma) \geq u_m(\neg\phi)$ for arbitrary m, which is just to say that $\Gamma \models_{sumdeg} \neg\phi$. So reductio is valid. This reasoning also shows that the converse of reductio holds.
5. **Transitivity** holds. Suppose $\Gamma \models_{sumdeg} \phi$ and $\phi \models_{sumdeg} \psi$. Then we know on any model, the sum undeterminacies of the premises of the first match or exceed the undeterminacy of ϕ, which matches or exceeds the undeterminacy of ψ. But then the transitivity of 'matching or exceeding' gives us: $\Gamma \models_{sumdeg} \psi$.
6. **Cut** fails. We can have cases where $\Gamma \models_{sumdeg} \phi$, $\Gamma, \phi \models_{sumdeg} \psi$, but $\Gamma \not\models_{sumdeg} \psi$. Illustrative examples of this were given earlier: for example p, $p \wedge D_{0.5} p \models_{sumdeg}$, $p \wedge D_{0.5} p \models_{sumdeg} p$ but $p \wedge D_{0.5} p \not\models_{sumdeg}$.
7. **Conditional proof** fails. Counterexample: $p, D_{0.5} p \models_{sumdeg} \neg p$. But $D_{0.5} p \not\models_{sumdeg} p \rightarrow \neg p$. To see the first holds, consider two cases. First, take a supervaluational model where the degree of determinacy of p is 0.5. Then the sum undeterminacy of the premises will be 0.5, and the sum undeterminacy of the conclusion will be 0.5 too. Second, take a supervaluational model where the degree of determinacy of p is not 0.5. Then $D_{0.5} p$ will have degree of determinacy 0, and the sum undeterminacy of the premises is at least 1, which must exceed the undeterminacy of the conclusion. Either way, the condition for degree-validity is met.
8. **Conjunction introduction** fails. Counterexample: We noted above that $p \wedge D_{0.5} p \models_{sumdeg} \neg p$ and clearly $p \wedge D_{0.5} p \models_{sumdeg} p$; but we also know that $p \wedge D_{0.5} p \not\models_{sumdeg} p \wedge \neg p$—the premise can have degree of determinacy 0.5 in a model, while the conclusion has always degree of determinacy 0. So we've a counterexample to the metarule that would take us from $\Gamma \models B$, $\Gamma \models C$, to $\Gamma \models B \wedge C$.[41]

These results are stated for $\models_{sumdeg}$—that is, what we have lately been calling $\models^0_{sumdeg}$. What of the formal properties of other degrees of validity? In fact, most of these results go through for these relations too. The argument for supraclassicality is given below. Reflexivity and monotonicity go through for the same reasons as above—for example, since sum undeterminacy can only increase when extra premises—and if the sum-undeterminacy+k exceeded the undeterminacy of the conclusion in the original case, it will do so with the extra premise added. Reductio holds too. Suppose $\Gamma, \phi \models^k_{sumdeg}$. That requires for arbitrary m, $\sum_{\gamma \in \Gamma} u_m(\gamma) + u_m(\phi) + k \geq 1$. This is equivalent to $\sum_{\gamma \in \Gamma} u_m(\gamma) + k \geq 1 - u_m(\phi)$. But $1 - u_m(\phi) = u_m(\neg\phi)$, and hence we have $\sum_{\gamma \in \Gamma} u_m(\gamma) + k \geq u_m(\neg\phi)$ for

[41] Note that the failure of this metarule form of conjunction introduction is compatible with the validity of the sequent form: $\Gamma, \phi, \psi \models_{sumdeg} \phi \wedge \psi$.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

arbitrary m, which is just to say that $\Gamma \models^{k}_{sumdeg} \neg\phi$. This argument also reverses, giving both reductio and converse reductio.

Cut fails for $\models^{k}_{sumdeg}$ for $k \in (0, 1)$ just as it does for the original case. The original case suffices for $k < 0.5$. For greater k, we need to change the example, choosing $D_x p$ for varying x. The main difference when $k > 0$ in the properties listed is transitivity. This holds for 'full validity' that is $\models^{0}_{sumdeg}$, as argued above. But it fails for $\Gamma, \phi \models^{k}_{sumdeg}$ where $k \in (0, 1)$.[42]

B Proof of supraclassicality of $\models_{sumdeg}$. We want to show that for sequents $\Gamma \models \psi$ with finite premise sets, where all formulae are D-free (i.e., are sentences in the language of classical propositional logic) the sequent is valid on a classical reading iff it is valid on a sum-degree reading.

From right to left, it will suffice to show that any classical countermodel can be turned into a sum-degree countermodel (i.e., classical invalidity entails sum-degree invalidity). Take a classical countermodel i. Now consider the supervaluational model m with a single delineation d, and such that the interpretation I maps a sentence letter A to T at d iff $i(A) = T$. The truth values of all (D-free) formulae in the language are the same under i as they are under I at d. Finally, the measure must give weight 1 to d. We know that each $\gamma \in \Gamma$ is true on i, so it is true at d, m. So $\mu_m(\gamma) = 1$ for each such Γ, and the sum undeterminacy is 0. ϕ is false on i, and hence false at d, m. So $\mu_m(\psi) = 0$, and hence the undeterminacy of ψ is 1. So sum undeterminacy increases over the sequent in m, making m a countermodel to the sum-degree validity of the sequent. QED.

From left to right, we build on work in probabilistic logic. First the single-premise case. Suppose $\phi \models \psi$ is a classically valid sequent. Then on all classical interpretations where ϕ is true, ψ is true. It follows that in any supervaluational model, whenever ϕ is true on a sharpening, ψ is true on that sharpening. Thus the set of sharpenings on which ϕ is true on m is a subset of the set of sharpenings on which ψ is true on m. Given that admissible supervaluational models must make all sentence letters, and hence all formulae, true at a measurable subset of the sharpenings of m, we can conclude that $\mu_m(X) \leq \mu_m(Y)$, where X is the set of sharpenings in m making ϕ true, and Y is the set of sharpenings making ψ true. It follows that the undeterminacy of ϕ on this model must match or exceed that of ψ, showing that the sequent is sum-degree valid.

For the multipremise case, we note that for finite sets of premises, $\gamma_1, \ldots, \gamma_n \models \psi$ iff $\gamma_1 \wedge \ldots \wedge \gamma_n \models \psi$ if both sequents are read classically. We want to show that whenever the first sequent is classically valid, it is also sum-degree valid. Given the first is classical valid, so is the second, and hence by the single-premise result, for all supervaluational models m, $u_m(\gamma_1 \wedge \ldots \wedge \gamma_n) \geq u_m(\psi)$. It suffices therefore to show that for all supervaluational models, $u_m(\gamma_1) + \ldots + u_m(\gamma_n) \geq u_m(\gamma_1 \wedge \ldots \wedge \gamma_n) \geq u_m(\psi)$.

This can be shown by induction on n. $n = 1$ is trivial, and for the induction step we need to show that $u_m(A) + u_m(B) \geq u_m(A \wedge B)$. But this follows straightforwardly from the fact that μ_m is a measure. Let the set of sharpenings making A true be X, and the set of sharpenings that make B true be Y. The set of sharpenings making $A \wedge B$ true is $X \cap Y$, and those making $A \vee B$ true are $X \cup Y$. By definition, $u_m(A) = 1 - \mu_m(X)$,

[42] Let's suppose we have 'constants' α_x for various degrees of determinacy x. For $k \in (0, 1)$, consider the sequents $\alpha_1 \models \alpha_{1-k}$ and $\alpha_{1-k} \models \alpha_{1-k-\epsilon}$. For $\epsilon < k$, these both will be k-valid, but $\alpha_1 \models \alpha_{1-k-\epsilon}$ will not be.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

$u_m(B) = 1 - \mu_m(Y)$, $u_m(A \wedge B) = 1 - \mu_m(X \cap Y)$, and $u_m(A \vee B) = \mu_m(X \cup Y)$. A defining feature of measures that are at least finitely additive is that $\mu(X) + \mu(Y) = \mu(X \cap Y) + \mu(X \cup Y)$. By the earlier equations, this is to say that $1 - u_m(A) + 1 - u_m(B) = 1 - u_m(A \wedge B) + 1 - u_m(A \vee B)$. Rearranging, $u_m(A) + u_m(B) = u_m(A \wedge B) + u_m(A \vee B)$. And hence, by the property that the measure never assigns a weight of less than zero, we have $u_m(A) + u_m(B) \geq u_m(A \wedge B)$. This is what was required for the inductive hypothesis, and so the result follows. QED.

What about arguments that are sum-degree valid to any degree other than 1? The first part of the above argument goes through directly. Any classical countermodel is a sum-degree-k countermodel (since what the above constructs is a supervaluational model where there is a measure 1 increase in undeterminacy). Moreover, since we've shown that any classically valid sequent is sum-degree-0 valid (i.e., no increase in undeterminacy allowed) it'll follow a fortiori that such classically valid sequents never involve an increase in undeterminacy of greater than k. So $\models^{k}_{sumdeg}$ are supraclassical in the stated sense for every $k \in [0, 1)$.

BIBLIOGRAPHY

Adams, E. (1996). *A Primer of Probability Logic*. Chicago, IL: CSLI.

Akiba, K. (2000). Vagueness as a modality. *Philosophical Quarterly*, **50**, 359–370.

Barnes, E. J., & Williams, J. R. G. (2010). A theory of metaphysical indeterminacy. *Oxford Studies in Metaphysics*, **6**.

Edgington, D. (1997). Vagueness by degrees. In Keefe, R. and Smith, P., editors. *Vagueness: A Reader*. Cambridge, MA: MIT Press, pp. 294–316.

Field, H. (2009). What is the normative role of logic? *Proceedings of the Aristotelian Society*, **83**, 251–268.

Field, H. H. (1973). Theory change and the indeterminacy of reference. *Journal of Philosophy*, **70**, 462–481. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001), pp. 177–198.

Field, H. H. (1974). Quine and the correspondence theory. *Philosophical Review*, **83**, 200–228. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001), pp. 199–221.

Field, H. H. (2000). Indeterminacy, degree of belief, and excluded middle. *Nous*, **34**, 1–30. Reprinted in Field, *Truth and the Absence of Fact* (Oxford University Press, 2001), pp. 278–311.

Field, H. H. (2003). Semantic paradoxes and the paradoxes of vagueness. In Beall, J. C., editor. *Liars and Heaps*. Oxford: Oxford University Press, pp. 262–311.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, **30**, 265–300. Reprinted with corrections in Keefe and Smith, editors. *Vagueness: A Reader* (Cambridge, MA: MIT Press, 1997), pp. 119–150.

Halpern, J. Y. (1995). *Reasoning About Uncertainty* (revised edition). Cambridge, MA: MIT Press. Revised paperback edition published 2003.

Joyce, J. M. (1998). A non-pragmatic vindication of probabilism. *Philosophy of Science*, **65**, 575–603.

Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Huber, F. and Schmidt-Petri, C., editors. *Degrees of Belief*. Berlin: Springer, pp. 263–297.

Kamp, J. A. W. (1975). Two theories about adjectives. In Keenan, E., editor. *Formal Semantics of Natural Language*. Cambridge, MA: Cambridge University Press, pp. 123–155. Reprinted in Davis and Gillon, editors. *Semantics: A Reader* (Oxford: Oxford University Press, 2004), pp.541–562.

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press

Keefe, R. (2000). *Theories of Vagueness*. Cambridge, MA: Cambridge University Press.
Lewis, D. K. (1970). General semantics. *Synthese*, **22**, 18–67. Reprinted with postscript in Lewis, *Philosophical Papers I* (Oxford University Press, 1983), pp. 189–229.
Lewis, D. K. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, **62**(3), 221–236. Reprinted in Lewis, *Papers on Metaphysics and Epistemology* (Cambridge University Press, 1999), pp. 56–77.
Lewis, D. K. (1993). Many, but almost one. In Campbell, K., Bacon, J., and Reinhardt, L., editors. *Ontology, Causality and Mind: Essays on the Philosophy of D. M. Armstrong*. Cambridge, MA: Cambridge University Press. Reprinted in Lewis, *Papers on Metaphysics and Epistemology* (Cambridge University Press, 1999), pp. 164–182.
Macfarlane, J. (2003). Future contingents and relative truth. *Philosophical Quarterly*, **53**, 321–336.
Macfarlane, J. (2010). Fuzzy epistemicism. In Moruzzi, S. and Dietz, R., editors. *Cuts and Clouds*. Oxford: Oxford University Press, pp. 438–463.
Machina, K. F. (1976). Truth, belief and vagueness. *Journal of Philosophical Logic*, **5**, 47–78. Reprinted in Keefe and Smith, editors. *Vagueness: A Reader* (MIT Press, 1997), pp. 174–204.
Paris, J. (2001). A note on the Dutch Book method. In *Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications, ISIPTA*, Ithaca, NY: Shaker, pp. 301–306.
Restall, G. (2005). Multiple conclusions. In Hajek, P., Valdes-Villanueva, L., and Westerstahl, D., editors. *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*. London: King's College Publications, pp. 189–205. Available from: http://consequently.org/writing/multipleconclusions/.
Smith, N. J. J. (2008). *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.
Smith, N. J. J. (2010). Degrees of truth, degrees of belief and subjective probabilities. In Moruzzi, S., and Dietz, R., editors. *Cuts and Clouds*. Oxford: Oxford University Press, pp. 491–506.
Stalnaker, R. (1980). A defense of conditional excluded middle. In Harper, R., Stalnaker, W. L., and Pearce, G., editors. *Ifs: Conditionals, Belief, Decision, Chance and Time*. Dordrecht, Holland: Kluwer Academic Publishers, pp. 87–106.
Thomason, R. H. (1970). Indeterminist time and truth-value gaps. *Theoria*, **3**, 264–281.
Varzi, A. (2007). Supervaluationism and its logics. *Mind*, **116**(463), 633–676.
Weatherson, B. (2003). From classical to constructive probability. *Notre Dame Journal of Formal Logic*, **44**, 111–123.
Williams, J. R. G. (2007). Eligibility and inscrutability. *Philosophical Review*, **116**(3), 361–399.
Williams, J. R. G. (2008). Supervaluations and logical revisionism. *The Journal of Philosophy*, **105**, 192–212.
Williams, J. R. G. (manuscript). Gradational accuracy and non-classical logic.
Williamson, T. (1994). *Vagueness*. London: Routledge.

DEPARTMENT OF PHILOSOPHY,
UNIVERSITY OF LEEDS
WOODHOUSE LANE
LEEDS, WEST YORKSHIRE, LS2 9JT, UNITED KINGDOM
Email: j.r.g.williams@leeds.ac.uk

https://doi.org/10.1017/S1755020310000237 Published online by Cambridge University Press