

Degree supervaluational logic

J. R. G. Williams

January 19, 2010

Abstract

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. 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 multi-premise 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. More worryingly, metarules such as Cut and conjunction introduction fail. It is argued that such failures do not make the logic 'damagingly revisionary'.¹

¹Special thanks are owed to: Elizabeth Barnes, Jc Beall, Rachael Briggs, Ross Cameron, John Divers, Daniel Elstein, Greg Restall, and especially Jason Turner. Thanks to all with whom I've discussed this material, including audiences at a CMM seminar in Leeds and at the 2009 AAP in Melbourne.

The work in this paper was funded by a British Academy Research Development Award (BARDA: 53286); and was supported by the workshops and materials provided by Spanish Government grant FFI2008-06153 (MICINN).

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, Achille Varzi 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. (Varzi, 2007, p.463)

The 'supervaluational machinery' to be discussed here is the idea of a *supervaluational model* defined below. Varzi highlights the fact that all sorts of properties of sequents that are candidates for the name 'validity' can be defined using the resources of supervaluational models.

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.

In this paper I focus on the role of logic in imposing (static) coherence norms on belief states, a key idea being that credence should not drop from the premise to conclusion of a valid argument. Rival characterizations of logic are then argued to impose different (and sometimes incompatible) coherence norms on belief states. In particular, they correspond to different views of the 'cognitive' or 'doxastic' role of determinacy judgements—what sort of attitude to p can legitimately be combined with high credence that p is indeterminate.

Hartry Field in recent work ((2003b), (2003a)) has argued for what a 'rejectionist' cognitive role for indeterminacy: believing that p is indeterminate rationally requires rejection of both p and of $\neg p$. Following Field (2003a), we shall see that one famous characterization of logical consequence using supervaluational machinery, construed normatively, enforces rejectionism.

But there are alternative takes on the cognitive role of indeterminacy. The main focus of this paper is on a ‘degreed’ cognitive role. This extends the thought that in paradigm cases where p is indeterminate, one’s credences in p should be 0.5. More generally, one’s credence in p should match the *degree of determinacy* of p —where paradigm cases of indeterminacy are supposed to have degree of determinacy 0.5. I associate this with a slight modification of the supervaluational framework proposed by David Lewis (1970), Hans Kamp (1975) and Dorothy Edgington (1997).² Thinking in terms of degrees of determinacy makes salient many more characterizations of ‘consequence’, but we shall see that most of these either undergenerate or overgenerate normative constraints on belief states.

The conclusion I argue for is that the logic of determinacy that norms belief in this setting is peculiar in two ways. First, in order to capture the full range of rational constraints, we need to talk not just in terms of a binary distinction between validity and invalidity, but in terms of degrees of validity (a notion characterized in a rather different context by Machina (1976)). Second, though the logic is non-revisionary in the sense that every classically valid argument remains valid (it is ‘supraclassical’) we find distinctly non-classical behaviour in a language that includes determinacy operators. This situation is familiar from the literature on global supervaluational consequence, which while supraclassical doesn’t sustain classical metarules such as conditional proof or reductio (Williamson, 1994, ch.5). I’ll show that degree-validity in fact retains reductio; but conditional proof still fails. In fact, we have even more dramatic departures: the structural rule of ‘Cut’, and the metarule of conjunction introduction both fail.

My view is that the ‘supervaluationist’ who thinks of the cognitive role of indeterminacy in a rejectionist way will favour a global characterization of consequence; one who thinks of the cognitive role of indeterminacy in the degree-norm way will favour the degree-logic characterization given here. In each case, there are marked departures from classical logic. But these are only two among the host of possible characterizations of consequence. The challenge

²All of these author’s treatments involve subtle variations that I ignore here. In the case of Lewis and Kamp, these are mostly focused on getting right the compositional clauses for various natural language terms (in particular, modifiers and comparative forms of adjectives). The variations in Edgington’s presentation are more directly concerned with the status of the framework. Edgington argues directly for ‘degrees of determinacy’ that work very like those below, and mentions the Lewis-Kamp framework as one articulation of this idea. But she doesn’t commit herself to it, identifying what she thinks of as additional commitments of the kind of ‘degree supervaluational’ model that she does not want to buy into.

to those favouring some rival definition of consequence is to describe an integrated package of (i) views on the cognitive role of indeterminacy; (ii) normative role for logic and (iii) characterization of logic, so we can see what advocacy of their favoured characterization as the correct ‘logic for determinacy’ comes to.

In **section 1** I characterize supervaluational models and give some illustrative applications. In **Section 2** I give a standard characterization of a supervaluational determinacy operator. I supplement the models with a measure over their ‘sharpenings’, and use this to characterize *degree determinacy* operators in similar fashion. **Section 3** gives two familiar characterizations of logic in terms of supervaluational model theory—local and global consequence. Via a connection between logic and rational belief, this is connected to the rejectionist conception of indeterminacy introduced above. **Section 4** introduces a rival to the rejectionist picture, based on the idea of matching one’s credence to degrees of determinacy. **Section 5** and **Section 6** consider a variety of characterizations of consequence that exploit the degreed setting, and evaluates their fit with the degree-cognitive role of the previous section. **Section 7** and **Section 8** extend this notion by introducing the notion of *degrees of validity* and considering extensions to the case of multi-premise sequents. **Section 9** establishes a variety of metatheoretic properties of the degree logic we end up with (a proof of the central property of supraclassicality is given in the appendix). As already flagged, some of these metatheoretic properties are rather strange. I argue in the conclusion that the strangeness of these metarules does not mean the logic is ‘damagingly’ revisionary of ordinary inferential practice.

1 Basics

A basic supervaluational model $m = \langle D, S, I \rangle$ for a propositional language L consists of a set of ‘delineations’ D , an interpretation function I that assigns T or F to each sentence in L relative to each delineation.³ The model contains a set of delineations that are special—the sharpenings S .

³The *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.

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 .

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

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 non-deterministic, multiple delineations (histories), corresponding to historically possible evolutions of the universe, and thereby picked out. What is supertrue is what is determined to be true by the laws+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 (think of some ‘growing block’ theories, for example).⁴

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 where 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,

⁴See e.g. Thomason (1970). This sort of model, with tweaks that leave unchanged the basic behaviour, underlies the recent discussion of future contingents and relativism (Macfarlane, 2003).

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.⁵

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’.⁶ 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 true, and a sharpening where this is false. Such borderline sentences will then be neither supertrue nor superfalse.⁷ 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); the problem of the many (Lewis, 1993).

In these 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 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.

⁵See e.g. Akiba (2000), Barnes & Williams (2010).

⁶Lewis (1984), Williams (2007).

⁷See e.g. 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.

It is standard allow the object-language under consideration to express claims about what is or what is not indeterminate. The standard device for achieving this to add to the propositional base language a sentential operator we will 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.⁸ 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.

If there are only finitely many sharpenings in a model, this is an idea that anyone friendly to supervaluational models can accommodate. The degrees of determinacy underlying comparative determinacy claims can be identified with the *proportions* of the sharpenings that make the relevant sentence true. Thus, a sentence will be 0.9 determinate if it is true on ninety per cent of the sharpenings—and 0.1 determinate if it is true on ten per cent of them.

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.⁹ What is the measure? If there are only finitely many sharpenings, it is simply the

⁸The benefits they hope differ among each other, and 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.

⁹ 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 ‘non-measurable’. If S is infinite, we can’t guarantee there won’t be non-measurable sets. This leads to the

count-measure, reporting the proportion of sharpenings that lie within a given set. So if $X \subset S_m$, $\mu_m(X)$ equals $|X|$ divided by $|S_m|$. In the infinite case, it is no longer explicitly definable in this way, but—in the most basic case—it is intended to reflect facts about the relative number of sharpenings in a given set.

One might be sceptical of whether there is any content to the notion of the ‘relative number’ of items in an arbitrary subset of a given infinite set. I’m personally agnostic on this issue—the mere fact that ‘relative number’ measure is not definable in terms of ratios of cardinalities does not show that the conception is contentless. But to commit oneself to the claim that there is a privileged ‘pure counting’ measure in the infinite case will be contentious. Nevertheless: the simplest form of the supervaluational degree of determinacy will be committed to such a notion.

If there is no infinite ‘pure counting’ measure, then there are still many interesting applications of the degree-determinacy formalism. In particular, local to a particular application, one might claim we can describe how an ‘intended measure’ could be determined—not reflecting proportions of sharpenings (even in the finite case) but some other quantity. (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 amongst 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. The construction of such measures is likely to be philosophically controversial.)

Controversial or not, the logical issues developed below are unaffected by these differences in

question of how we handle the logic and semantics of sentence letters that are true exactly on a non-measurable set of sharpenings (cf. Smith (2008, pp.)). 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 will be measurable). The proof of supraclassicality below appeals to this.

A different salient option is 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 and μ can be explicitly defined—if we wish—as the pure counting measure. See the discussion below.

the ‘metaphysics of the measure’, so I set the issue aside. For definiteness, I’ll continue to talk as if it was a pure counting conception that is at issue.

If such a measure is to hand, we can make the following sort of distinctions in our sample applications. Harry is a paradigm borderline bald—it’s totally indeterminate whether or not he’s bald. But ten years ago, when he had a little more hair, the indeterminacy was a little less—it was slightly more determinate that he was hairy then, than it is now. This claim is vindicated if the relative proportion of sharpenings on which Harry is bald have increased over the years. Or consider a biased coin, which comes down heads 99 per cent of the time—it’s more determinate that it will come down heads than that a fair coin will (even though both, qua future contingents, are indeterminate). Again, the claim is vindicated if the biased coin comes down heads on more of the histories than does the fair coin.

To allow for such talk in our object language we need to add degreed determinacy operators. 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 .

In short, we are allowed talk of degrees of determinacy, corresponding (on the count-interpretation of the measure) to the proportion of sharpenings which make the statement true. We 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 exactly measure k of the sharpenings.

With this notion around, fans of the above semantic model can make sense of much of the talk about indeterminacy that ‘fuzzy logicians’ go in for.¹⁰ 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 non-baldness (perfectly determinate cases of non-baldness)—without any ‘big

¹⁰See Machina (1976) and Smith (2008).

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 non-zero degree of determinacy. Similarly for 'Harry is bald, and Jimmy is not' (where Jimmy is balder than Harry).

If one wishes, one can use the ideas just introduced to give precise sense to the idea of 'degrees of truth', characterized as follows:¹¹

- 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

But if one is uncomfortable about this talk (and uncertain about its relation to truth talk more generally) then one needn't go in for it. One could alternatively talk in the metalanguage of 'degrees of determinate truth' of sentences, or dispense with such metalinguistic glosses altogether, just talking directly of proportions of sharpenings making the sentence true.

3 Logic

There has been considerable debate about how friends of supervaluational models can and should 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

¹¹There 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 degree-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'.

‘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.

One tradition seeks to resolve this dispute by first answering a question which we have been silent about to this point—what is truth, in this setting? Given an answer to this question, and ‘platitudes’ connecting truth to logic, we may then get a fixed on what the right characterization of consequence is.¹² A traditional supervaluationist answer (*the* supervaluationist answer, on one understanding of ‘supervaluationist’) is that truth is supertruth. So sentences that are neither supertrue nor superfals are genuinely counterexamples to bivalence. Next, appeal to the ‘platitude’: logical consequence is guaranteed truth preservation. We get a characterization of logical consequence as ‘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

¹²See Williamson (1994, ch.5) for a version of this strategy.

¹³For now I focus on a single-premise, single-conclusion case since it is simpler to state, 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.

¹⁴Many 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 non-classical behaviour 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 defence of this as the right logic for D , see the final chapters of Keefe (2000)).

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

The ‘truth to logic’ reasoning above suggests that if we’re persuaded that we should identify truth (simpliciter) with supertruth, then we should accordingly think that \models_g , rather than \models_l , accurately reports the facts about what follows from what.¹⁵

Maybe the platitudes are worth preserving. We could then reduce the task of arguing over the correct logic, to arguing over the correct characterization of truth (which of course is a substantial issue in itself). But I think we can do better than this in understanding what’s at stake in a debate over ‘real logical consequence’. To determine what the correct characterization of logic is, I propose we look at what role logic plays for us: and then think of logic as whatever plays the logic-role.

There are numerous ways in which logic interacts with wider theory, but the role I will focus on here is to do with *norming belief states*. Start with the thought that 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 are subject too a range of conflicting evidence. But that shouldn’t blind us to the fact that inconsistency is a serious (not merely prima facie) flaw in a belief state. Pro tanto problems with a belief state may be outweighed by competing considerations, but they remain in place. Here, then, is one logic-belief constraint: pro tanto, one’s belief state shouldn’t be inconsistent.¹⁶

If one favours a description of belief states in terms of *partial belief* or *credence*, rather than all-or-nothing belief, then the logic-belief norm just sketched won’t have direct application.

But there’s a compelling formulation for this case (which, happily for us, shifts the focus from inconsistency to validity, and so is directly applicable to the kind of simple formulations of

¹⁵This also famously forms the basis for the notorious failure of classical metarules for \models_g (Williamson, 1994, ch.5). For even though $p \models_g Dp$, $\not\models_g p \supset Dp$. 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.

¹⁶Maybe the constraint kicks in only when the inconsistency is or should be obvious—I leave this wrinkle to one side for now.

validity we've been formulating): pro tanto, one shouldn't have less credence in the conclusion of a valid sequent than in its premise.

If we insist that when asking about 'logic', we are at least asking for something that plays the belief-state norming role, then there is real bite to debates about what logic the friend of supervenient models should endorse. Suppose I'm fully confident that Harry is borderline bald. I'm credence zero that it's determinate that he's bald. But as noted above, this is a g -consequence of Harry 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—zero. Mutatis mutandis for 'Harry is not bald', and we get the result that I should be credence zero in both 'Harry is bald', and its negation. 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. ¹⁷

This gives life to the debate about how to characterize logic in the supervenient setting. I take it that the friend of rejectionism will favour 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's the alternative to the global/rejectionist package?

¹⁷A formal model of credences appropriate to rejectionism (suggested by Field) can be found in the literature on Dempster-Shafer belief functions. See Halpern (1995).

One way of capturing these 'pro tanto belief norms' systematically is to formulate a (possibly non-classical) *probability theory*—together with the claim that ideal credences will satisfy those axioms. Classical probability theory can be axiomatized via the classical consequence relation—and we get a generalized notion of logic by switching out classical consequence for some non-classical consequence. See Weatherson (2003) for discussion of this for the case of intuitionistic logic and an accompanying notion of 'constructive probability'. In (Williams, manuscript) I discuss these generalizations in a very general setting, and adapt the 'accuracy' arguments of Joyce (1998) to tie together model theory, logic and probability in the generalized setting.

4 Degree cognitive role

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.

One might dislike this particular answer. One natural thought is that faced with a smack-bang borderline case of baldness (like Harry) your credences should be 50/50. 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 supervenient 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. 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 the biased coin mentioned earlier will land heads. On a natural articulation of the view, this matches up with the degrees of determinacy in each case.

So the 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.¹⁸ This gives a distinct

¹⁸This is a familiar thought in the literature on non-supervenient degrees of truth—see Smith (forthcoming) and Macfarlane (forthcoming).

cognitive role to indeterminacy judgements. If you are certain that the degree of determinacy of p is 0.5, then you know what your first-order opinion should be: you should have 0.5 credence in p .

Without a take on the cognitive role of indeterminacy judgements, we're in a weak position to discriminate between candidate characterizations of 'logic' in the belief-norm sense. Rejecting rejectionism requires taking a (negative) stance on such issues. But now we've seen an alternative, positive view. The task in the remainder of the paper is to explore the kind of position on logic that this view motivates.

5 Characterizations of logic by degrees

\models_g is not an appropriate logic for one favouring the degree-cognitive role just described, since it enforces rejectionism. There's no such objection to the local characterization. But so far, we've only consider a couple of candidate characterizations.¹⁹ In the present context it's natural at least to look to characterizations in degree-theoretic terms.

There is a literature that's interestingly analogous to the present discussion: 'probability logic'. This studies—among other things—constraints on how probabilities interact with (classically) valid sequences. And it turns out there are numerous ways of giving *purely probabilistic* characterizations of logic for standard languages, which extensionally coincide with classical validity.²⁰

'Degrees of determinacy' and 'probabilities' are formally identical, and so we can simply take over the work that is done in probability logic, to arrive at candidate formulations of a 'degree-determinacy logic':

¹⁹For several more, see Varzi (2007). Some of these come apart only when multiple premises and conclusion are in play. But a salient alternative to the local and global options is the 'backwards falsity preservation'—which corresponds to a cognitive role where p and $\neg p$ should both be accepted when p is indeterminate. This characterization fits most naturally with a 'subvaluationist' gloss on which truth is identified with truth on some sharpening, and so indeterminate sentences get to be both true and false. I won't consider this further here.

²⁰See Adams (1996).

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

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

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

At this point, we can piggyback on technical results proved in probability logic to show that all the above definitions coincide with classical logic on the D -free fragment of the language we are working with.²¹

The above characterizations coincide with each other, and with local and global validity, for the D -free fragment of the language. But we know from the global/local validity debates, that tweaking the characterization of validity can have dramatic effects that show up only when determinacy operators are introduced.

So it is here. Here are four results illustrating extensional divergence between our various characterizations:

- $p \models_1 D_1 p$. But $p \not\models_b D_1 p$ and $p \not\models_{deg} D_1 p$
- $p \models_b D_{\geq 0.999} p$. But $p \not\models_{deg} D_1 p$
- $D_1 p \models_{deg} p$

The first result is that on the ‘strict’ or ‘1-preservation’ characterization, but on neither of the alternatives, that p is determinate to degree 1 will follow from p itself. To see this, note that if p is degree 1 determinate in m , then by the above definitions, $D_1 p$ will be true on each sharpening of m , and hence determinate to degree 1 as well. On the other hand, $D_1 p$ will be determinate to degree 0 whenever p falls short of degree 1 determinacy—no matter how small the shortfall is. And that means this isn’t valid on the bounds characterization. This also shows that on the degree-preservation characterization the sequent is invalid.

²¹For the case of degree-preservation, see the theorem of supraclassicality in the final sections of this paper. For others, see the proofs in Adams (1996). The appendices of that book also sketch how the approach can be generalized to a first-order language, which will be a useful model for developing the degree determinacy approach to that setting.

The second result is that, on the bounds characterization, it follows from p that p is determinate to a very high degree—since we can force the latter to itself be determinate to degree 1 by making p determinate enough. (It is pretty obvious that 1-preservation also supports this sequent). But the degree-preservation characterization doesn't vindicate this sequent—if p is true on half the sharpenings and false on half, then it'll be determinate to an intermediate degree but $D_{\geq 0.999}p$ will be determinate to degree zero—a fall in determinacy that provides a counterexample to the sequent.

Together, these two initial results show that the three degree-invoking characterizations of validity diverge from one other in a language that includes determinacy operators.

We can rerun much of the discussion over global consequence with the 1-preservation and bound characterizations of validity. In particular, the logics will not sustain classical metarules. Take conditional proof, for example. Despite the above results, we don't have $\models_1 p \supset D_1p$. Take a model on which p is true on only half the sharpenings. Then D_1p is false on all the sharpenings. The material conditional in question is therefore false on exactly those sharpenings on which p is true. So it is determinate to degree 0.5. So it's not a 1-tautology.

Likewise, we don't have $\models_b p \supset D_{\geq 0.999}p$. Take the same model as above—again, the determinacy claim is false on each sharpenings, and so the conditional is false on those sharpenings where p is true. So the conditional is again determinate to degree 0.5. This shows it is not a bound-tautology. So conditional proof fails for 1-preservation and bound-logic.

As well as the failure of these metarules, we can also see that construing these logics as norming belief states would lead to rejectionism. Suppose I am certain that p is determinate to degree 0.5. I'm then in a state where I'm credence 0 that p is determinate to degree at least 0.999. So by credence-preservation across valid arguments, we'd be committed by a normative construal of the bound-logic to assigning credence zero to p (and equally, to $\neg p$). Exactly similar considerations show that 1-preservation enforces rejection, if it is treated as a norm for belief states.

If we want to avoid rejectionism, we can rule out global consequence, bound consequence, and

1-preservation consequence as logics that norm belief.

6 Local and Degree-preservation logic

Let's now consider the degree-preservation logic defined above. We've already seen ways that it diverges from bound and 1-preservation logics. As a side effect, it's not immediately obvious that metarules like conditional proof will fail. In fact, they do—but I'll delay discussion of the various metatheoretic properties of this notion until the final sections (suffice to say that it's behaviour is *very* unlike classical logic—even a metarule like 'the conjunction of two logical consequences of A is itself a consequence of A ' lapses).

Let's turn to the final result given above: that $D_1p \models_{deg} p$ (recall that D_1p is shorthand for $D_{\geq 1}p$ —it is true (on a model) exactly when measure one of the sharpenings of that model are true. This sequent holds on the degree-preservation logic characterized above. Indeed, for D_1p to be determinate to any positive degree in m , p must hold on measure-one of the sharpenings of m . So there's no possibility of a drop in the determinacy-degree from premise to conclusion, just as required. (It is also valid on the 1-preservation and bound logics, but they're out of the running by this stage).

However, it's arguable that the sequent isn't valid on the local characterization of consequence. It might be that D_1p is true on m (i.e. p is true on measure 1 of the sharpenings of m), without p being true on every sharpening of m .²² To get a feel for how this might happen, consider the following famous puzzle case: if one flips a fair coin infinitely many times, what's the probability of it coming up heads each time? Well, the possibility of it coming up heads on each of the first n occasions is $1/2^n$, and it coming up heads on every occasion is surely strictly smaller than that. But that means that the all-heads result must be smaller than $1/2^n$ as n tends to infinity, and on standard assumptions this forces the probability of all heads to be zero (one can try to avoid this reasoning by appealing to 'infinitesimal' probabilities—I won't consider

²²This can happen only when there are infinitely many sharpenings in S . As emphasized earlier, when there are only finitely many sharpenings in S the 'count measure' is explicitly definable in the metalanguage, and so this is the one place in our discussion where the appeal to measures within the model is irreducible.

this here). The same goes, in fact, for every particular infinite tails-heads sequence. This shows that we just can't think of 'probability zero' events as impossible—else each completed infinite sequence would be impossible! Every particular outcome of an infinite sequence of heads is probability zero—and so, conversely, there is probability 1 of the outcome being other than all-heads (say). Probability is just a certain kind of measure, so this illustrates the following properties of measures: measure 1 of the outcomes in a space can have a particular property (e.g. being other than all-heads) while some of the outcomes in the space lack the property (in particular, the all-heads outcome).

Suppose we have this situation with determinacy measures: measure 1 of the sharpenings make p true, but at least one amongst them makes p false. Those sharpenings will then constitute a countermodel to the claim that $D_1 p \models_l p$. For in that model, on that sharpening, the antecedent is true and the consequent false. To avoid this result one would have to argue that, for some reason, there are no legitimate supervaluational models where this situation arises. I see no reason to think this would be the case. So I conclude that the sequent given above is locally invalid.

I said that it's a mark in favour of a logic that norms belief (on the conception of the cognitive role of indeterminacy) if it captures the fact that being fully confident of $D_1 p$ rationally requires being fully confident of p . So the determinacy-degree logic earns a good mark, and the local validity characterization fails to do so.

The problems for local validity here are different from those faced by global, 1-preservation and bound validity—for those predicted a constraint that conflicts with the non-rejectionist conception of the cognitive role we're investigating. Here, the problem is one of undergeneration rather than overgeneration—local validity fails to predict something that *does* appear to be a rationality requirement on belief states. A similar situation would arise if we asked whether local validity is an adequate logic for the rejectionist—the worry would not be that it imposes extra rationality requirements that conflicts with what the rejectionist wants to endorse, but that it fails to capture the distinctive constraints on belief which the rejectionist believes determinacy-judgements impose.

We shall see in the final section of this paper that degree-logic, though well motivated, has peculiar metatheoretic properties. One in particular can be appreciated even in the single-premise setting that we have dealt up to now. It turns out that a version of *conjunction introduction* fails. In particular, we will have instances where $A \models_{deg} B$ and $A \models_{deg} C$, but $A \not\models_{deg} B \wedge C$. This will be shown, and its significance evaluated, at the end of the paper.

Having flagged up this issue, in the next two sections I consider extensions of the notion of degree logic. To fully capture norms for belief, even the degree-preservation logic undergenerates. However the degree-characterization, unlike local validity, has a natural generalization that plausibly does spell out the relevant rationality requirements. In the section that follows, I consider generalizations to multi-premise sequents.

7 Degrees of validity

\models_{deg} correctly articulates a rationality requirement connecting belief involving a determinacy operator ($D_1 p$) to beliefs in the embedded claim (p). But it doesn't do all one might hope.

Consider, for example, the following sequent:

$$p \wedge D_{0.5} p \models$$

On a normative construal of the logic involved, if correct this would mean that we'd be pro tanto wrong in assigning the conjunction more than credence zero.

But this isn't a valid sequent in the degree-determinacy logic (nor in the local logic). And a good thing too. On the conception of indeterminacy we're after, $p \wedge D_{0.5} p$ can be determinate to degree 0.5 (on a model where exactly half the sharpenings make p true). On the conception of indeterminacy we're working with, your credence in the conjunction should be 0.5. We certainly don't want our logical norms to enforce credence 0 here!

On the other hand, there are some rationality constraints between p and $D_{0.5} p$. It'd be

irrational for you to be certain of the latter, and not assign 0.5 credence to p , for example. Or, considering the conjunction itself, it's noteworthy that the maximal indeterminacy you can have in it is 0.5. Unless it is determinate to exactly this degree, it'll be perfectly 0-determinate (to see this, just suppose that the case mentioned above doesn't obtain—that the measure of sharpenings in m on which p is true is not 0.5. In that case, $D_{0.5}p$ is false on each sharpening, and hence the conjunction is false on each sharpening, and hence determinate to degree zero). So these kind of conjunctions are very special—they're something it'd be irrational ever to fully believe. But there's no way of extracting this constraint on belief states out of the characterization of the logic above. We're again undergenerating. (The friend of a local characterization of validity might reasonably ask: why worry if we miss a rationality constraint linking credences in D_1p and p , if we're anyway going to have to give an independent story linking p and $D_{0.5}p$?)

There is a natural generalization of the determinacy-degree story which can plug the gap here.²³ In particular, let's see degree validity as the limiting case of a whole family of degreed notions of validity. To repeat the earlier characterization in new notation:

- $\phi \models_{deg}^0 \psi$ 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:

- $\phi \models_{deg}^k \psi$ 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 .

Notice that though $p \wedge D_{0.5}p \models_{deg}^0$ doesn't hold for the reasons give above, $p \wedge D_{0.5}p \models_{deg}^{0.5}$ does. For by the above, 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.

²³The notion is prefigured, in quite a different setting, by Machina (1976).

Our original logic-credence link was this: pro tanto, one's credences shouldn't drop over a valid argument. Spelling this out in terms of degree-validity, that's to say: pro tanto, one's credences shouldn't drop over an argument whose degree of determinacy is guaranteed not to drop. And this suggests a natural generalization: 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 . That is: 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.

This generalization 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.

Let me emphasize a couple of things in closing this section. First, I see no way of replicating this kind of generalization for the local setting. So while degree-determinacy logic generalizes to a package that plausibly captures the distinctively logical constraints on beliefs about determinacy, local validity does not. The charge that local logic undergenerates the rationality constraints is sustained. Second, I've been careful to present all the above in terms of one role for logic: norming belief states. This makes the degreed validity relation interesting even if for other purposes they're badly behaved. For example, drops in determinacy can 'add up' over the course of several arguments. This leaves open that $A \models_{deg}^{0.9} B$ and $B \models_{deg}^{0.9} C$ but $A \not\models_{deg}^{0.9} C$. So transitivity may fail. For a logic that hoped to capture good patterns of inference, that's arguably a pretty devastating blow—how are we to theorize about sustained inference if we can't chain together valid arguments? However, it's no criticism at all of bringing in non-transitive notions to do the job I've prescribed for them.

Finally, the whole discussion thus far has a serious limitation. We've described the normative role of logic, but have only characterized validity for single-premise arguments.

Correspondingly, while we can say what it is for a single formula to be inconsistent (or indeed to have various degrees of inconsistency— $p \wedge D_{0.5}p$ being naturally described as inconsistent to degree 0.5, for example), we haven't yet got the resources to say what it is for a set of claims to be inconsistent. It is to this we now turn.

8 Multiple premises

How might we generalize the ideas of the degree-preservation to the deal with multiple premises?²⁴

One prominent idea can be drawn from the literature on probabilistic logic.²⁵ First, let's define the degree of *undeterminacy* (lack of determinacy) of p in a model, as the measure of the sharpenings on which p is false in that model. Let's write this as $u_m(p)$. The original idea of validity as a guarantee that determinacy won't drop over an argument can be expressed in these terms as a guarantee that the undeterminacy won't increase.

The basic idea of the 'summing' conception of multi-premise validity is that consequence relation is a matter of a guarantee that the *total* undeterminacy doesn't increase. So the sum total undeterminacy of the premises should be greater than or equal to the undeterminacy of the conclusion:

summing degree preservation

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

It is a famous result in probability logic that this kind of summing characterization coincides with classical multi-premise validity *at least for a language involving no probability or determinacy operators*.²⁶

²⁴For full generality, we might also wish to generalize to the case of multiple conclusions—see Varzi (2007), Restall (2005). I'm not going to consider this extension here, however.

²⁵Adams (1996)—where this is used as the definition of 'probabilistic validity'.

²⁶See 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.

One could of course, consider other extensions. If one is willing to trivialize the equivalence between $A \wedge B$ on the one hand, and A, B taken individually, then one could simply reduce the multi-premise case to the single-premise case:

conjunctive degree preservation

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

This of course would only work (in a finitary language) when there are only finitely many premises. But we can reasonably stick with this for now. In this setting, in a language without D -operators these two characterizations coincide (this result is provable using elementary features of classical propositional logic, and the results from probability logic mentioned earlier).

But again, we should look for divergences in the setting where we have determinacy operators in our language. And we get very surprising results in this setting. In particular, the summing conception leads to failures of *idempotency*—the structural claim that one can eliminate repeated premises from a valid inference without losing validity. The summing degree logic is *substructural*.²⁷

To see a case of this, consider the set of premises $\{p, D_{0.5}p\}$. When p is determinate to degree other than 0.5, this sum total undeterminacy will be at least 1, since the other claim will be determinate to degree 0. Suppose on the other hand that p is determinate to degree 0.5. Then $D_{0.5}p$ has a degree of undeterminacy of 0, and p has a degree of determinacy of 0.5, so the sum total is 0.5.

So just as $p \wedge D_{0.5}p \not\models_{deg}$, we have

$$p, D_{0.5}p \not\models_{deg}$$

On the other hand, consider the triple $\{p, p, D_{0.5}p\}$.²⁸ Again, where p has degree of

²⁷Thanks to Jason Turner for pointing this out to me.

²⁸To get this case, we need to deal with multi-sets—sets where the number of repetitions of an element matters to identity.

indeterminacy other than 0.5, the sum total indeterminacy will be at least 1, due to the last element. But where p is determinate to degree 0.5, it counts *twice*, and so the sum total indeterminacy is again at least 1. Since this means that any claim whatsoever will be no greater than the sum total indeterminacy of the set, we have:

$$p, p, D_{0.5}p \models_{sumdeg}$$

Here we have the failure of idempotency—the latter displayed sequent is inconsistent, but it merely repeats a premise already present in the earlier displayed sequent.

One might be tempted to try tweaking the definition (say, by identifying the validity of a sequent with repeated premises with the validity of the sequent where repetitions are eliminated—making idempotency trivial). But the non-classicality here runs deep. Contrast, for example the inconsistency of $\{p, p \wedge D_{0.5}p\}$ and the consistency of $\{p \wedge D_{0.5}p\}$. This occurs despite the fact that $p \wedge D_{0.5}p \models p$ on any of the logics we have seen. Here, without repeated premises, we have the result that adding logical consequences of a premise set as an additional premise can move one from consistency to inconsistency—or equivalently, that eliminating premises which are already logical consequences of what remains can turn a valid argument invalid. Both this and the earlier idempotency-failure illustrate cases where $\Gamma, \phi \models \psi$ and $\Gamma \models \phi$, but where $\Gamma \not\models \psi$. That is, we have a failure of the ‘Cut’ rule.

The conjunctive extension to the logic is much more conservative. It will be idempotent and support Cut. For $\phi, \phi, \Gamma \models_{conjdeg} \psi$ is just to say $\phi \wedge \phi \wedge \bigwedge_{\gamma \in \Gamma} \gamma \models_{conjdeg} \psi$. But clearly $\phi \wedge \phi \wedge \bigwedge_{\gamma \in \Gamma} \gamma \models_{conjdeg} \phi \wedge \bigwedge_{\gamma \in \Gamma} \gamma$ and hence by transitivity of single premise sequents:²⁹ $\phi \wedge \bigwedge_{\gamma \in \Gamma} \gamma \models_{conjdeg} \psi$, which then means $\phi, \phi, \Gamma \models_{conjdeg} \psi$.

Let’s think in more detail, however, about whether the conjunctive characterization really will do the job we want it to. How should we characterize the constraints that facts about multi-premise validity put on sets of beliefs (ignoring for a moment the issue of

²⁹for proof of transitivity see below

generalizations to degrees of validity)?

It's easiest to think about this in terms of inconsistency (where X is inconsistent iff every sentence in the language is a consequence of it). In normal cases, if a set of sentences Γ is inconsistent, then this is by no means 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. Moreover, Γ 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.

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. The standard way to articulate this constraint will look familiar. If we write $\bar{c}(p)$ for to represent the degree to which you're uncertain of p (i.e. 1-your credence in p), then the pro tanto constraint on credences when $\Gamma \models$ is:

$$\sum_{\gamma \in \Gamma} \bar{c}(\gamma) \geq 1$$

And accordingly, if we're consider a valid argument where $\Gamma \models \phi$, the standard line is that the accompanying rationality constraint is that:

$$\sum_{\gamma \in \Gamma} \bar{c}(\gamma) \geq \bar{c}(\phi)$$

Where there is a single premise, then this comes down to the constraint that $\bar{c}(\gamma) \geq \bar{c}(\phi)$, which is equivalent to the constraint that your credence in γ should be no more than your credence in ϕ —so we have as a special instance of the above the single-premise norm that we've been working with thus far.

Let's think of this in connection to the above examples. A useful case study will be $\{p, p \wedge D_{0.5}p\}$. I contend that here there should be a constraint that one's uncertainties in the

two claims mentioned should add up to at least 1. For, as noted before, one's credence in $p \wedge D_{0.5}p$ shouldn't ever be more than 0.5. And for your credence in the conjunction to be even this high, you must have 0.5 credence in p . If we turn to the pair above, your degrees of uncertainty will therefore add up to more than 1.

$\{p, p \wedge D_{0.5}p\}$ is inconsistent on the summing characterization, but isn't on the conjunction characterization. And the logic-belief link formulated above will then, correctly, impose the constraint that uncertainties should add up to more than 1. But the set is not inconsistent on the 'conjunction characterization' of multi-premise validity, and hence it misses this constraint. Moreover, the difference in status between $\{p, p \wedge D_{0.5}p\}$ (inconsistent on the summing conception) and $\{p \wedge D_{0.5}p\}$ (consistent) is a precondition of generating the rationality constraints here. For while our uncertainties must add up to over 1 in the former, there is no constraint that they must do so in the latter case. Since the conjunction formulation collapses the distinction between the former and the latter, there's little prospect of it generating the correct constraints. The conjunctive characterization undergenerates rational constraints on credences.

We know, from before, that to get the full range of rationality constraints, we'll have to relate differences in credence to degrees of validity. We therefore need to generalize the notion of degreed validity, and formulate the required constraints. In each case it's obvious how to take the first step:

summing degree preservation

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

And there's a natural corresponding generalization of the logic-belief state constraint given earlier: whenever $\Gamma \models_{sumdeg}^k \psi$, it's pro tanto irrational to have one's sum total uncertainty in the premises be less than the uncertainty one has in the conclusion *by more than k*.

9 Summary of the properties of sumdeg validity

I have been arguing that someone with a ‘degreed’ conception of the cognitive role of determinacy judgements should think that the logic that norms belief-states is the degree-logic characterized above. Before wrapping up, I want to summarize its metatheoretic properties, proving various results that have been alluded to above (in one case, I defer the proof to an appendix).

1. It is **Supraclassical** for (finite) multi-premise sequents. That is, whenever $\Gamma \models \phi$ is classically valid (where Γ is finite and all formulae are D -free), we have $\Gamma \models_{sumdeg} \phi$. The proof of this is given in an appendix.
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 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 $\Gamma \models B \wedge C$.

The final result is particularly strong, since it only involves single-premises sequents. This means that the counterexample to the metarule-formulation of conjunction introduction goes through even in the *single-premise* fragment of degree-logic \models_{deg} . So *whatever* stance we take on the generalization to multi-premise sequents, we’re not going to avoid this kind of strangeness. (Note that this metarule isn’t the only thing that might go by the name ‘conjunction introduction’. I have not suggested a countermodel to the single sequent $A, B \models A \wedge B$, for example).

These results are stated for \models_{sumdeg} —i.e. what we have lately been calling \models_{sumdeg}^0 . 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 in an appendix. 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 $\Gamma, \phi \models_{sumdeg}^k$. 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 arbitrary m , which is just to say that $\Gamma \models_{sumdeg}^k \neg\phi$. This argument also reverses, giving both reductio and converse reductio.

Cut fails for \models_{sumdeg}^k 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’ i.e. \models_{sumdeg}^0 , as argued above. But it fails for $\Gamma, \phi \models_{sumdeg}^k$ where $k \in (0, 1)$.³⁰

³⁰Let’s suppose we have ‘constants’ for various degrees of determinacy α_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 sufficiently small ϵ , these will both be k -valid, but $\alpha_1 \models \alpha_{1-k-\epsilon}$ will not be.

10 Conclusion

Let's consider where we've come from. We started with a certain kind of semantic model—supervaluationist models, incorporating many sharpenings. Appeal to such things are widespread in philosophy—in the philosophy of time, in paradigmatic vague predicates, in recent work on metaphysical indeterminacy.

But these models famously leave open the question of what the correct 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. As urged by Hartry Field in particular, one kind of doxastic role is that spelled out by the rejectionist: judging p to be indeterminate requires that one should reject p and also reject $\neg p$. One would hope that these kind of normative principles would flow from a pair of claims (i) about what the logic of indeterminacy is; (ii) about the normative role of logic. Global consequence, together with a standard 'no drop in credence' construal of the normative significance of logical consequence, plays this role for the rejectionist.

But what if we have some rival idea about the cognitive role that indeterminacy should play? Here I've been investigating a non-rejectionist conception on which one's strength of belief in p should (ideally) match the proportion of sharpenings on which p is true (whether we want to introduce the ideology of 'degrees of truth' at this point doesn't really matter—we certainly should have some way of expressing these proportions in the object-language, but I've been working with 'degrees of determinacy' to try and emphasize the continuity with 'supervaluational' rather than 'degree theoretic' thoughts). The question arises: what sort of characterization of logic should we give, if we want it to play the normative role just specified? Many natural candidates overgenerate—claiming rational constraints that simply are not there on the conception of indeterminacy in question. Others undergenerate. We're left with just one characterization standing—the one I've labelled summing degree-preservation.

In the literature on supervaluationism, it is often emphasized that global consequence, while

Note to self—find replacements for constants that exist within the official language to make this point.

matching classical logic over what counts as a logical truth, departs from classical logic at the level of metarules. What goes for global consequence goes in spades for summing degree consequence. While global consequence lacked certain metarules (once we have degreed determinacy operators in the object language) summing degree consequence lacks classical structural metarules such as Cut. That may be good or bad for the overall view, but I've tried to emphasize that it's needed in a logic that wants to capture exactly the rationality requirements that this conception of determinacy imposes (the contrast with the conjunctive characterization makes this point).

Of course, our whole discussion has proceeded on the assumption that the 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.³¹ 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 these extended 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.

Let me finish with one last suggestion. In Williams (2008), it is noted that the failure of classical metarules doesn't preclude the obtaining of *weakened versions* in the context of 'global supervaluational consequence'.³² 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

³¹It'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?

³²A similar point is made by Vann McGee, and in recent work by John Macfarlane.

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 goodstanding are suddenly declared unreliable, so long as they fit the pattern of the weakened versions of the metarules.

Degree logic, and sum degree logic are the best motivated characterizations of consequence for the degree supervaluationist. The metatheory is strange, as I've outlined. But I don't think the sort of revisionism involved undermines familiar, classical-seeming *inferential practice*; and so I do not see this revisionism as the source of an *objection* to characterizing logic in the way indicated.

A Appendix: 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 multi-premise case, we note that for finite sets of premises, $\gamma_1, \dots, \gamma_n \models \psi$ iff $\gamma_1 \wedge \dots \wedge \gamma_n \models \psi$ 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 , $u_m(\gamma_1 \wedge \dots \wedge \gamma_n) \geq u_m(\Psi)$. It suffices therefore to show that for all supervaluational models, $u_m(\gamma_1) + \dots + u_m(\gamma_n) \geq u_m(\gamma_1 \wedge \dots \wedge \gamma_n) \geq u_m(\Psi)$.

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, $u_m(A) = 1 - \mu_m(X)$, $u_m(B) = 1 - \mu_m(Y)$, $u_m(A \wedge B) = 1 - \mu_m(X \cap Y)$ 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 a increase in undeterminacy of greater than k . So \models_{sumdeg}^k are supraclassical in the stated sense for every $k \in [0, 1)$.

References

- ADAMS, ERNST. 1996. *A primer of probability logic*. Chicago, IL: CSLI.
- AKIBA, KEN. 2000. 'Vagueness as a modality'. *Philosophical Quarterly*, **50**, 359–370.
- BARNES, E. J., & WILLIAMS, J. ROBERT G. 2010. 'A theory of metaphysical indeterminacy'. *Oxford Studies in Metaphysics*, **6**.
- 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. 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, HARTRY H. 2003a. 'No fact of the matter'. *Australasian Journal of Philosophy*, **81**, 457–480.
- FIELD, HARTRY H. 2003b. 'Semantic paradoxes and the paradoxes of vagueness'. *Pages 262–311 of: BEALL, J. C. (ed), Liars and Heaps*.
- 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.
- HALPERN, J. Y. 1995. *Reasoning about uncertainty*. Revised edn. MIT Press. Revised paperback edition published 2003.
- JOYCE, JAMES M. 1998. 'A non-pragmatic vindication of probabilism'. *Philosophy of Science*, **65**, 575–603.
- 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.
- LEWIS, DAVID K. 1970. 'General Semantics'. *Synthese*, **22**, 18–67. Reprinted with postscript in Lewis, *Philosophical Papers I* (Oxford University Press, 1983) 189–229.
- LEWIS, DAVID K. 1984. 'Putnam's paradox'. *Australasian Journal of Philosophy*, **62**(3), 221–36. Reprinted in Lewis, *Papers on Metaphysics and Epistemology* (Cambridge University Press, 1999) 56–77.
- LEWIS, DAVID K. 1993. 'Many, but almost one'. In: CAMPBELL, KEITH, BACON, JOHN, & REINHARDT, LLOYD (eds), *Ontology, Causality and Mind: Essays on the philosophy of D. M. Armstrong*. Cambridge: Cambridge University Press. Reprinted in Lewis, *Papers on Metaphysics and Epistemology* (Cambridge University Press, 1999) 164–82.

- MACFARLANE, JOHN. 2003. 'Future contingents and relative truth'. *Philosophical Quarterly*, **53**, 321–336.
- MACFARLANE, JOHN. forthcoming. 'Fuzzy epistemicism'. In: MORUZZI, SEBASTIANO, & DIETZ, RICHARD (eds), *Cuts and Clouds*. Oxford: OUP.
- 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.
- RESTALL, GREG. 2005. 'Multiple Conclusions'. *Pages 189–205 of: HAJEK, PETR, VALDES-VILLANUEVA, LUIS, & WESTERSTAHL, DAG (eds), Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*. London: King's College Publications. Available online at <http://consequently.org/writing/multipleconclusions/>.
- SMITH, NICK J. J. 2008. *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- SMITH, NICK J. J. forthcoming. 'Degrees of truth, degrees of belief and subjective probabilities'. In: MORUZZI, SEBASTIANO, & DIETZ, RICHARD (eds), *Cuts and Clouds*. Oxford: OUP.
- STALNAKER, R. 1980. 'A defense of conditional excluded middle'. In: HARPER, STALNAKER, & PEARCE (eds), *Ifs: Conditionals, Belief, Decision, Chance and Time*. The Netherlands: Kluwer Academic Publishers.
- THOMASON, RICHMOND H. 1970. 'Indeterminist time and truth-value gaps'. *Theoria*, **3**, 264–281.
- 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.
- WILLIAMS, J. ROBERT G. 2007. 'Eligibility and inscrutability'. *Philosophical Review*, **116**(3), 361–399.
- WILLIAMS, J. ROBERT G. 2008. 'Supervaluations and logical revisionism'. *The Journal of Philosophy*, **CV**(4).
- WILLIAMS, J. ROBERT G. manuscript. 'Gradational accuracy and non-classical logic'.
- WILLIAMSON, TIMOTHY. 1994. *Vagueness*. London: Routledge.