Normative indeterminacy in the epistemic domain

Nick Leonard and Fabrizio Cariani

1 Introduction

This paper is about epistemic conflicts in which one is doomed to violate a rule of rationality no matter what. Here is one example, lifted from Leonard (2020):

> LOGIC PROBLEM: Anna is a logic student who is evaluating the following proposition:
>
> (L) If Lenny is happy if and only if Jenny and Benny are not happy, then if Jenny is happy if Benny is not, Lenny is not happy.
>
> Anna is certain that (L) is a tautology, and is thus true. However, her logic professor tells her that before she began the exam, she was slipped a reason distorting drug that impairs one's ability to solve logic problems; those who are affected by the drug only reach the right conclusions 20% of the time. As it turns out, unbeknownst to Anna, the drug was a placebo and Anna's logic reasoning abilities were not affected in the least.[1]

What should Anna believe? We are pulled in two directions by reflecting on different aspects of epistemic rationality. On the one hand, consider

Probabilism: An agent's credences ought to be probabilistically coherent.[2]

Probabilism demands that Anna's credence in (L) be 1.

On the other hand, consider,

---

[1] See Christensen (2007, 2010) for similar cases.

[2] Influential defenses of Probabilism include Skyrms (1966), De Finetti (1992), Christensen (1996), Joyce (1998), and Pettigrew (2016).

Calibrationism: For any proposition, **p**, if a rational agent forms a credence in **p** on the basis of some first-order reasoning, and if, independently of the first-order reasoning in question, that agent's expected degree of reliability concerning whether **p** is r, then the agent's credence in **p** ought to be r.[3]

Anna should expect her degree of reliability concerning (L) to be .2. By Calibrationism, then, Anna's credence in (L) should be .2

Thus, agents like Anna are subject to epistemic rules that cannot be jointly satisfied, opening up the question what epistemic rationality requires of them.

Such situations of conflict need not be this far-fetched. Here is a different case lifted, with some modifications, from Knoks (2020):

> IMPLICIT BIAS. Suppose that there are two kinds of people, circles and squares. Sammy is a square. She's been exposed to evidence that very strongly suggests that almost all squares harbor an implicit bias which causes them to believe, incorrectly, that circles are far ruder than they really are. Sammy meets a circle, Carl, and judges them rude on the basis of several personal interactions. Unbeknownst to Sammy, though, she is unusual in that she doesn't suffer from this particular implicit bias.

Sammy's interactions with Carl provide her with great evidence that they are rude. However, she also has higher-order evidence to the effect that she has likely misjudged the evidence regarding Carl's impoliteness.

Sammy faces a conflict between two other seemingly plausible epistemic principles:

Evidential Support: If an agent's evidence strongly supports believing **p**, then they ought to believe **p**.[4]

Interlevel Coherence: A rational agent ought to be such that (i) if they believe that their evidence supports believing **p**, then they ought to believe **p**

---

[3]Influential defenses of Calibrationism include Roush (2009),White (2009), Sliwa and Horowitz (2015), and Christensen (2016). While Calibrationists disagree about how exactly this principle should be formulated, the consensus is that Anna's credence in (L) should be .2.

[4]Some (e.g., Stapleford and McCain (2020)) reject Evidential Support in favor of a weaker principle according to which you are only required to believe what your evidence strongly supports if you have actually taken an attitude towards **p**. The conflict can be captured in terms of this principle too, since we can stipulate that Sammy has considered whether Carl is rude.

and (ii) if they believe that their evidence does not support believing **p**, then they ought not believe **p**.

Evidential Support entails that Sammy should believe that Carl is rude. However, because she has good reason to think that she suffers from implicit biases, this principle also entails that Sammy should believe that the evidence does not support her belief that Carl is rude.[5] On the basis of this, Interlevel Coherence requires her to not believe that Carl is rude.[6] Like Anna, then, Sammy is subject to normative epistemic constraints that cannot be jointly satisfied.

The general question behind these puzzles is: what does epistemic rationality require us to think when its constraints are jointly unsatisfiable?[7] Our goal in this paper is to develop an answer to this question according to which, in situations of conflict, epistemic rationality is indeterminate. When one is subject to conflicting epistemic rules, it is indeterminate which rules one is permitted and required to satisfy, and thus indeterminate which doxastic attitudes one is permitted and required to have. To meet our goal, we are going to argue for two main claims. First, that rational indeterminacy can be modelled within the framework of defeasible deontic logic, as proposed by Horty (2003, 2012) and then applied to epistemic conflicts by Knoks (2020a,b). We show that the detour through this logical framework is essential to spelling out the view in full generality. Second, that this approach provides a unified way of treating epistemic conflicts while also enjoying some important advantages over its rivals.

Here is the plan. In section 2, we informally characterize three different approaches to epistemic conflicts. In section 3, we model each approach in our preferred deontic logic — a slight variant of Knoks's system. In sections 4 and 5, we offer some new reasons for preferring our view to its rivals. Finally, in section 6, we compare our way of modelling rational indeterminacy to a

---

[5]We are assuming that Sammy has good reason to think that her implicit biases cause her to be so wrong in her initial assessment of Carl's rudeness that she should think that her evidence does not actually support the claim that Carl is rude.

[6]Similar cases are discussed in Littlejohn (2018), Worsnip (2018) and Lasonen-Aarnio (2020).

[7]Other epistemic conflicts are discussed in Conee (1987), Sorensen (1987), Egan and Elga (2005), Caie (2012, 2013), and Schechter (2013).

simpler supervaluationist treatment developed in Leonard (2020).

2   Three approaches to epistemic dilemmas

In this section we provide a brief, informal overview of three possible answers to our central question. The approaches agree that in situations of conflict rational agents are not doomed to violate one of their epistemic requirements.[8] At this stage — prior to the technical development of section 3 — we focus on the general motivation of each proposal, rather than on any precise statement.

We start with a permissivist position.

Choice. When there are multiple epistemic rules that cannot be jointly satisfied, one is permitted but not required to satisfy any maximally satisfiable combination of them.[9]

In other words, rationality is permissive in the sense that one is free to adopt any sufficiently strong combination of rules to satisfy, as long as they satisfy one of them.

According to the second view, the rules of rationality are structured by priority relations. When the rules conflict, one is required to satisfy a subset of them that includes the highest ranking ones, if there are any, and perhaps some others.

Priority. Rules do not have equal weight. When multiple rules cannot be jointly satisfied, priority relations between them contribute to the determination of the final epistemic oughts.[10]

For instance, if the demands of one's first order evidence conflict with the demands of one's higher-order evidence, and if Evidential Support is given more weight than Inter-Level Coherence, then one is required to believe what their first-order evidence supports. Depending on the priority structure, there might not be a highest ranking option, in which case one would determine the final epistemic oughts on the basis of those rules that are not defeated by

---

[8]For a fourth position according to which one can be subject to epistemic requirements that cannot be jointly satisfied, see, e.g., Priest (2002), Ross (2010), Brouwer (2014), Hughes (2019), and, on one interpretation, Christensen (2007, 2010). And see Horty (2012) for one way of formally modeling this approach. Finally, see, Caie (2012), Turri (2012), Lasonen-Aarnio (2014), and Leonard (2020) for some recent objections to this view. Because adequately engaging with this position would take us beyond the scope of this paper, we will set it aside.

[9]See, e.g., (Knoks, 2020b, §5).

[10]See, e.g., (Knoks, 2020b, §6) and, on one interpretation, Christensen (2007, 2010).

anything else. It is worth noting that in this kind of case, one might consider blending a priority based picture with a permissivist one.

Lastly, there is the option that conflicting epistemic rules give rise to normative indeterminacy.

Indeterminacy. When there are multiple epistemic rules that cannot be jointly satisfied, it is indeterminate which maximally consistent combination one is permitted or required to satisfy. One is, however, determinately required to satisfy the disjunction of maximally consistent combinations.[11]

Thus if two rules are in conflict, it is indeterminate whether one is permitted and required to satisfy the first, and it is indeterminate whether one is permitted and required to satisfy the second. Nevertheless, one is determinately required to satisfy one or the other.

It is important to appreciate some subtleties when comparing Indeterminacy to the alternatives. Indeterminacy shares with Choice the idea that one is determinately required to satisfy one of the conflicting rule-sets. Indeterminacy, however, stops short of declaring it *permissible* to satisfy any one combination. Furthermore, the normative indeterminacy position is compatible with the central contention of the Priority-based theorist, i.e., that rules can be assigned various weights.

With these informal characterizations in hand, we are now going to highlight how each one can be formally modeled.

3   Formal analysis

Our preferred model is a relatively small reinterpretation and modification of Knoks' (2020b) formalism. Exploring the details of Knoks's system serves two purposes. First, the model offers an explanation of how epistemic requirements and permissions are generated from the rules of rationality. For instnace, instead of treating Evidential Support and Interlevel Coherence as indefeasible rules that one is always required to satisfy, this model treats them as defeasible rules that contribute to the final determination of what rational agents are epistemically permitted and required to think, i.e., these defeasible rules ground the selection of binding epistemic principles, but they need not be identical with them. The second reason for exploring this model is that it frames one of our main contributions—the formal modeling of a sophisticated indeterminacy-friendly position.

---

[11]See, e.g., Leonard (2020).

3.1 The language

At the beginning of our analysis, we consider a language with the following primitive grammar:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \bigcirc\varphi \mid \boxed{b}\varphi \mid !(\varphi \mid \varphi) \mid \varphi \Rightarrow \varphi,$$

Informally, this adds three operators to a basic sentential language: '$\bigcirc$' is an operator tracking the epistemic requirement the agent is bound by; '$\boxed{b}$' is a belief operator. Next, '!' is a binary operator that expresses conditional rules. In the defeasible frameworks of Horty (2003, 2012), which Knoks builds on, requirements and permissions emerge from *rules*.[12] The process involves a series of devices whose job is to individuate which rules are to prevail, once we take into account their compatibility and priority relations. The syntax of '!' is vaguely reminiscent of the syntax of a conditional probability operator. In particular, in the expression '$!(\varphi \mid \psi)$, ψ is a condition under which the rule with content φ is triggered. Lastly, '$\Rightarrow$' denotes a support relation, holding between a body of evidence (modeled as a set of propositons) and an individual propositions. Although we generally follow Knoks's system quite closely, our construal of the support relation is a bit different from his. In Knoks's system, the relation denoted by '$\Rightarrow$' links a proposition representing the evidence and an attitude (e.g. a belief). We treat the evidence relation as relating the evidence and a proposition. Informally, instead of saying that E supports believing (the proposition expressed by) φ, we say that E supports φ. This makes no difference to the formal analysis but makes our explanatory task a bit simpler. Syntactically, $\Rightarrow$ is like a standard conditional, an unlike our '!' operator in having its first argument be the 'antecedent', so to speak, of the support relation.

3.2 Modeling

The first step in deploying this language consists in giving precise representations of the core epistemic rules from (Section 1) above. In IMPLICIT BIAS, for instance, Evidential Support amounts to:

ES. $!(\boxed{b}\varphi \mid E \Rightarrow \varphi)$

---

[12] Horty calls these *defaults* and Knoks calls these *imperatives*. We choose "rules" because we judge it to be more descriptive than these options.

Formalizing the Interlevel Coherence rule involves facing an important choice point. It is tempting and perhaps most elegant to formalize it as:

- $!(\neg\,\boxed{\text{b}}\,\varphi \mid \boxed{\text{b}}\,(E \nRightarrow \varphi))$

Informally: given the hypothesis that one believes that the evidence does not support φ, there is reason to not believe φ. Instead of this formalization, Knoks opts for fomralizing this as a rule with conditional content but no condition of its own. Formally, this is represented by the rule having a tautological ($\top$) hypothesis, as in:

IL. $!(\boxed{\text{b}}\,(E \nRightarrow \varphi) \supset \neg\,\boxed{\text{b}}\,\varphi \mid \top)$

This says that there is unconditional ('standing' in Knoks's terminology) reason to not believe those propositions that one believes are not supported by the evidence. The difference is subtle, but has important ramifications in the modeling of the cases. Once again, we follow Knoks.

Sticking with IMPLICIT BIAS, let p be an atomic sentence of the object language representing the proposition *Carl is rude*. Sammy's evidential situation can be represented as:

E1. $E \Rightarrow p$

E2. $E \Rightarrow (E \nRightarrow p)$

The evidence supports p and it supports the claim that the evidence does not support p. Each of these specific evidential claims is matched by an important instance of ES—one of the first-order proposition p and one for the higher-order support claim $E \nRightarrow p$.

ESF. $!(\boxed{\text{b}}\,p \mid E \Rightarrow p)$

ESH. $!(\boxed{\text{b}}\,E \nRightarrow p \mid E \Rightarrow (E \nRightarrow p))$

Similarly instantiating **IL** yields:

IRI. $!(\boxed{\text{b}}\,(E \nRightarrow p) \supset \neg\,\boxed{\text{b}}\,p \mid \top)$

This completes the representation of the principles involved in IMPLICIT BIAS.

3.3 Reasoning with the model

Knoks's modeling work illuminates the problem of normative conflict at the heart of epistemic dilemmas. In IMPLICIT BIAS, for instance, each of **ESF** and **ESH** is conditional on some evidential condition which in fact holds in the given situation. So **ESF** gives the agent reason to believe p and **ESH** gives her reason to believe $E \nRightarrow p$. An agent who enters both these states will find herself in conflict with **IRI**. And this brings us back to our key question: what are agents required and permtted to believe in situations like these?

Our final task in this section is to show how Priority, Choice, and Indeterminacy can be implemented in the default logic framework. The idea is that each position is associated with a different way of generating epistemic requirements and permissions from the epistemic rules that one is subject to in a given context. Contexts here are understood as abstract representations of the parameters that are relevant to the generation of epistemic requirements from epistemic rules. We may think that each belief state is held against the background of some information, as well as some salient epistemic principles. Thus understood, this concept of context is foundationally distinct from otherwise similar concepts of contexts philosophers of language might appeal to.

Choice

The permissivist twist on the system — the Choice option — is a valuable starting point. At this stage, a context c is modeled as a pair $\langle W, R \rangle$, where W is a set of formulas representing the factual circumstances of a given situation, while R is a set of rules. (Remember that rules are statements of the form $!(\psi \mid \varphi)$.) Given a rule, it is helpful to pick out its antecedent and consequent. For example, if $r = !(\psi \mid \varphi)$, $Antecedent(r) = \varphi$ and $Consequent(r) = \psi$. These functions are usefully polymorphic so that if R is a set of rules, $Consequent(R)$ is the set of consequents of the rules in R.

In our sample model of IMPLICIT BIAS, the context is pinned down by letting $W_1 = \{E1, E2\}$, while $R_1 = \{ESF, ESH, IRI\}$. We make this official by defining:

Context definition 1 $c_1 = \langle W_1, R_1 \rangle$.

We proceed to give a series of definitions that, given an arbitrary context $\langle W, R \rangle$ output the epistemic requirements that the agent is subject to in that

context. First, we identify the set of rules (from R) whose conditions are met in the situation:

Triggered(c) = $\{r \in R \mid W \vdash Antecedent(r)\}$[13]

A bit more precisely, but still informally, this definition distinguishes those rules in R whose precondition is established by the factual information in W from those whose precondition is not so established. To illustrate, imagine a context in which one rule is to believe what Tweedledee says, and another rule is to believe what Tweedledum says. Now imagine that only Tweedledee has spoken, then the former rule but not the latter rule is triggered.

Let us turn our focus back on c_1, our model of IMPLICIT BIAS. In this artfully chosen scenario everything is triggered. That is:

$$Triggered(c_1) = R_1 = \{ESF, ESH, IRI\}$$

To see this, note that ESF is triggered by $E1$, ESH is triggered by $E2$ and IRI is always triggered, because its precondition is a tautology. The set of conclusions of all of these triggered rules is the inconsistent set:

$$\{\boxed{b}(p), \boxed{b}(E \nRightarrow p), [\boxed{b}(E \nRightarrow p) \supset \neg \boxed{b}(p)]\}$$

The permissivist analysis generates epistemic oughts and permissions from this set: φ is epistemically obligatory iff it follows from all of the subsets of this set that are maximally consistent with W.

More formally, we provide the following chain of definitions.

Supported(c) = $\{\varphi \in \mathcal{L} \mid \varphi \in Consequent(Triggered(c))\}$

Maxicons$_X$(Y) = $\{Z \subseteq Y \mid X \cup Z \text{ is consistent and } \forall Z' \supset Z, X \cup Z' \text{ is inconsistent}\}$

Requirements(c) = $\{\varphi \in \mathcal{L} \mid \forall X \in MaxiCons_W(Supported(c)), W \cup X \vdash \varphi\}$

In order: the *supported* sentences are the conclusions of triggered rules. Our definition of maximal consistency is standard, with the exception of the fact that it is relativized to a set of formulas X. Finally, and crucially, the epistemic requirements in c are the maximally consistent (with respect to W) sets of supported formulas. It is possible to think of the last definition as implicitly

---

[13]In this and the subsequent definitions, we implicitly assume that W and R on the right side come from the context c on the left side, unless otherwise specified.

providing an analysis of an obligation operator $\bigcirc$ that draws on epistemic requirements. The analysis is that $\bigcirc\varphi$ holds in c iff φ is a requirement in c.

Let us return to our example to see why this epistemic ought captures the permissivist position. Note that the maximally consistent subsets of $Triggered(c_1)$ are:

(i) $\boxed{b}(p)$, $\boxed{b}(E \Rightarrow p)$

(ii) $\boxed{b}(p)$, $\boxed{b}(E \Rightarrow p) \supset \neg\boxed{b}(p)$ (via modus tollens this reduces to $\neg\boxed{b}(E \Rightarrow p)$)

(iii) $\boxed{b}(E \Rightarrow p)$, $\boxed{b}(E \Rightarrow p) \supset \neg\boxed{b}(p)$ (via modus ponens this reduces to $\neg\boxed{b}(p)$)

These correspond to the following conjunctions (we'll call these "takes" on the problem):

T1 $\boxed{b}(p) \wedge \boxed{b}(E \Rightarrow p)$

T2 $\boxed{b}(p) \wedge \neg\boxed{b}(E \Rightarrow p)$

T3 $\neg\boxed{b}(p) \wedge \boxed{b}(E \Rightarrow p)$

T1 corresponds to the fractured believer who disregards Interlevel Coherence. **T2** and **T3** respectively disregard the second order evidence or the first order evidence. According to the permissivist, none of the three takes on the problem is epistemically required, but (as Knoks clarifies), their disjunction is.

- $\neg\bigcirc(\mathbf{T1}) \wedge \neg\bigcirc(\mathbf{T2}) \wedge \neg\bigcirc(\mathbf{T3})$

- $\bigcirc(\mathbf{T1} \vee \mathbf{T2} \vee \mathbf{T3})$

Turning to permissions, recall that while Choice says that an agent is required to be in a state that satisfies the disjunction of the three takes, they are permitted to be in any of the three states. So, suppose we introduce a permission operator, with the stipulation that $\mathcal{P}\varphi$ holds at c iff $\neg\bigcirc\neg\varphi$ holds at c. Then all of the following hold at c_1:

- $\mathcal{P}[\boxed{b}(p) \wedge \boxed{b}(E \Rightarrow p)]$

- $\mathcal{P}[\boxed{b}(p) \wedge \neg\boxed{b}(E \Rightarrow p)]$

- $\mathcal{P}[\neg\boxed{b}(p) \wedge \boxed{b}(E \Rightarrow p)]$

This completes our formal presentation of Choice.

Priority

Knoks (2020b) also shows the intervention that is required to model the priority-based position. To take this step, we must enrich our conception of contexts. Instead of taking contexts to be pairs consisting of factual information and rules, we take them to be triples, by adding a priority component. The idea is that we do not just list epistemic rules, but also articulate their priority relations (again in the style of Horty (2003, 2012)). As a result, in this more sophisticated system, a context is a triple $\langle W, R, \prec \rangle$, where $\prec$ is a reflexive and transitive priority relation over rules. Intuitively $r_1 \prec r_2$ iff r_2 takes priority over r_1. Note that the relation may be partial, so that there may be incomparable rules.

Pinning down the epistemic requirements in these prioritized contexts is a more complex task. First, identify the consistent rules in c — those rules whose conclusions are compatible with the factual coordinate of c.

Consistent(c) $= \{r \mid Conclusion(r) \cap W \text{ is consistent }\}$

Next up, we identify the distinguished set of rules that are both triggered and consistent.

Active(c) $= \{r \mid r \in Triggered(c) \cap Consistent(c)\}$

Now let ρ be a variable ranging over sets of rules. We say that some rule r and a set of rules ρ are conflicting in c if they are all active in c and have incompatible conclusions.

Conflicted(c, ρ) $= \{r \mid Active(\rho) \wedge Consequent(\rho) \vdash \neg Consequent(r)\}$

For reasons related to our later development, it is convenient to present this as a property of a rule r that is relativized to a set of rules ρ.

Finally, Knoks borrows from Horty (2003, 2012) the idea of extending $\prec$ to relate sets of rules, as well as individual rules. In particular, $\rho > r$ means that every rule in ρ has priority over r.

We can now state the central concept of the priority-based analysis. A context may have several active rules, and those active rules may be in conflcit with each other. However, only a specified subset of those rules will be binding — that is, capable of generating epistemic requirements. These are those rules that are not *defeated* by any set of rules. That is to say those rules r such that there is no set ρ of conflicting rules that is uniformly of higher priority than r.

Binding(c) $= \{r \mid \neg \exists \rho \subseteq R, Conflicted(c, \rho) \wedge \rho > r\}$

With this concept in hand, we complete the analysis by replacing $Triggered(c)$ with $Binding(c)$ in characterizing the epistemic requirements.

Requirements(c) $= \{\varphi \mid \forall X \subseteq MaxiCons(Consequent(Binding(c))), W \cup X \vdash \varphi\}$

The priority analysis, as reconstructed by Knoks, yields different recommendations in IMPLICIT BIAS, depending on what kind of priority ordering is assumed in the context.

To start, consider a context c_2 that is like c_1 in its first two coordinates ($W_1 = \{E1, E2\}$, $R_1 = \{ESF, ESH, IRI\}$), with the additional information that the priority ordering $>_1$ is: $ESF >_1 ESH >_1 IRI$.

Context definition 2 $c_2 = \langle W_1, R_1, \prec_1 \rangle$

In c_2, any one rule from R_1 is conflicted relative to the pair consisting of the other two rules, but not relative to either rule individually. However, ESF and ESH are both binding. In the case of ESF, it is binding because nothing at all takes priority over it. As for ESH, it is binding because, although ESF takes priority over it, ESF isn't conflicting with it. Only the set consisting of ESF and IRI conflicts with ESH, and that set is not uniformly prioritized over ESH. As a result the requirements in c_2 are the consequences of the conclusions of ESF and ESH—i.e. the consequences of $\boxed{\text{b}}(p)$ and $\boxed{\text{b}}(E \nRightarrow p)$.

We note here a slight problem for Knoks's reconstruction of the priority system and suggest a solution. The concept of a binding rule relies on the idea of a set ρ of rules being better than a single rule r. As noted, Knoks obtains this relation by letting $\rho > r$ iff every member of ρ is uniformly prioritized over r. This is problematic in some cases. Consider for example the context c_3 defined as follows. Let the first two coordinates be as in c_1 and c_2; next define a partial order $\prec_2$ with $ESF \prec_2 ESH$ and $IRI \prec_2 ESH$, but with ESF and IRI incomparable with respect to each other.

Context definition 3 $c_3 = \langle W_1, R_2, \prec_2 \rangle$

The problem is that, under the current analysis, all three of ESH, ESF and IRI are binding in c_3. It is clear that ESH should be binding, since it has priority over both of the other rules. But now consider ESF: ESF is not defeated by ESH alone since it is compatible with it. It could potentially be defeated by the set $\{ESH, IRI\}$, but that set contains an element, IRI, that is not better than

ESF, so *ESF* is not defeated by any set that is uniformly better than it. For that reason the set of binding rules in c_3 is the entire set $\{ESF, ESH, IRI\}$. The upshot is that the ordering $\prec_2$ produces the same predictions as the ordering according to which all three rules are incomparable. The fact that *ESH* has priority over both *ESF* and *IRI* is washed out in the calculation. This shows up in the prediction that the system makes. For instance, it makes the implausible prediction that the agent is permitted to violate *ESH*.

A simple fix would be to modify how we extend the priority relations from individual rules to relations between sets of rules.[14] In particular, when ρ is a set and r is an individual rule, we can say that $\rho \succ r$ iff for every $s \in \rho$, $s \succeq r$ and for some $s \in \rho$, $s \succ r$. The effect of this change on the example involving c_3 is apparent and significant. As a result of it, the only binding rule in c_3 is *ESH*, and the epistemic requirements our agent is under arise by considering what follows from $\{ESH\}$. In sum, the agent ought to believe that the evidence doesn't support p but is under no obligation with respect to rule enjoining them to believe p itself or the higher-order constraint $\boxed{b}(E \Rrightarrow p) \supset \neg \boxed{b} p$.

Indeterminacy

The indeterminacy approach can be viewed as a generalization of the priority approach to a situation in which the context constrains, but does not fix, a priority structure.

Now, there are a few strategically diverging ways of presenting an indeterminacy based position. Consider again c_3. One basic thought is that this context determines that there are two sets of requirements that are not ruled out as binding. One set is $\{ESH, ESF\}$ and the other set is $\{ESH, IRI\}$. The priority ordering leaves indeterminate how *ESF* and *IRI* are to be ranked against each other.

We think of sets like $\{ESH, ESF\}$ as ways of resolving the indeterminacy in the priority ordering. To use a single word, we think of them as *resolutions* and we are going to construct them in two stages. First, we collect the maximally consistent sets of conclusions of binding rules: we call these the "proto-resolutions". Then, we form all the maximally consistent sets of conclusions of *triggered* rules that are compatible with some *proto-resolution*.

---

[14]There are a couple more complicated fixes one could think about. First, one might conceive of a system in which priorities have individual numerical weight, and then there is a module of the theory whose job is to extend that numerical weight from individual rules to sets of rules. Another option—one that stays within the qualitative framework of Knoks's analysis—is to adopt the replacement-based analysis in (Horty, 2012, ch.8).

Proto-resolutions(c) $= \{X \mid X \in MaxiCons_W(Consequent(Binding(c)))\}$

Resolutions(c) $= \{X \mid \exists Y \in Proto-resolutions(c), Y \subseteq X$
$$\text{and } X \in MaxiCons_W(Consequent(Triggered(c)))\}$$

So in our example, there is one proto-resolution, namely $\{ESH\}$. However, corresponding to it, there are two resolutions: $\{ESH, ESF\}$ and $\{ESH, IRI\}$. The distinction between proto-resolutions and full resolutions gives our theory a richer set of predictions which we won't explore until the last section of this paper.

Next up, we relativize requirements and permissions to a resolution.

Requirements$_X$(c) $= \{\varphi \mid W \cup X \vdash \varphi\}$

Permitted$_X$(c) $= \{\varphi \mid \neg\varphi \notin Requirements_X(c)\}$

Technically, of course, we have only relativized to a set of formulas X but in practice we think of these as resolutions of the indeterminacy.

Ultimately, it is desirable to have an unrelativized concept. We do this by characterizing the *determinate requirements*—those sentences which all of the individual resolutions uniformly agree about. These are obtained by closing off the relativization by quantifying universally.

Determinate Requirements(c) $=$
$$\{\varphi \mid \forall X \in Resolutions(c), \varphi \in Requirements_X(c)\}$$

Determinate Permissions(c) $=$
$$\{\varphi \mid \forall X \in Resolutions(c), \neg\varphi \notin Requirements_X(c)\}$$

More generally we say that given a relativized condition $\Theta_X(c)$, we have that $Det(\Theta(c))$ iff $\forall X \in Resolutions(c), \Theta_X(c)$.

Those sentences that are required (/permitted) relative to some, but not all resolutions, are indeterminately required (/permitted).

Indeterminate Requirements(c) $= \{\varphi \mid \exists X, Y \in Resolutions(c),$
$$\varphi \in Requirements_X(c) \wedge \varphi \notin Requirements_Y(c)\}$$

Indeterminate Permissions(c) $= \{\varphi \mid \exists X, Y \in Resolutions(c),$
$$\varphi \in Permitted_X(c) \wedge \varphi \notin Permitted_Y(c)\}$$

The indeterminist twist on the formalism is sensitive to differences in priority orderings. In a context like c_2 it will simply agree with Knoks's analysis. But its true force shines through when we consider a context in which the three rules have equal strength. In this kind of context, Indeterminacy makes a different recommendation from Choice, and also from Priority (which in this case just agrees with Choice).

Go back to our three takes from the analysis of the permissivist position:

T1 $\boxed{b}(p) \wedge \boxed{b}(E \nRightarrow p)$

T2 $\boxed{b}(p) \wedge \neg\boxed{b}(E \nRightarrow p)$

T3 $\neg\boxed{b}(p) \wedge \boxed{b}(E \nRightarrow p)$

Naturally, what turns out to be required on this approach depends significantly on the priority ordering.

Start by considering an ordering such that the three rules have equal priority. Then all three rules ESF, ESH and IRI are binding. Accordingly, we have three resolutions $\{ESF, ESH\}$, $\{ESF, IRI\}$ and $\{ESH, IRI\}$—respectively supporting one each of **T1, T2**, and **T3**. So for each one of our three takes it's indeterminate whether the agent ought to be in that state, but it is determinate that one ought to be in a state which satisfies the disjunction of **T1, T2**, and **T3**. While the view agrees with Choice on this score, it denies that one is determinately permitted to be in a state satisfying **T1** (and similarly for **T2**, and **T3**.)

Because we have added the indeterminacy module on top of a priority-based system, we are also able to coherently model those cases in which there *is* a priority ranking. Thus, consider again the hypothetical context with $ESF > ESH > IRI$. In this context, the only binding resolution is $\{ESF, ESH\}$. Thus, insofar as this really is how these rules should be ranked, it is determinate that one ought to satisfy **T1**.

4 Objections to Choice

The permissivist option, Choice, rests on an extreme form of intrapersonal permissivism. It can be permissible for one agent, with one relatively rich body of evidence and a single evidential standard (i.e., a way of weighing various pieces of evidence against each other), who is subject to one set of moral and practical stakes, to adopt one of a huge swath of credence levels in a particular proposition at a particular time.

In LOGIC PROBLEM, for instance, Choice says that even if everything is held fixed (e.g., stakes, evidence, Anna's evidential standards, time, etc.), Anna is still permitted to be certain that (L) is true, and she is also permitted to be 80 percent certain that it is false.

Insofar as rationality has something to do with proportioning your doxastic attitudes to your evidence, Extreme Intrapersonal Permissivism is a tough pill to swallow. Once everything is held fixed, the more Anna's evidence and evidential standards point towards (L) being true, the more they should point towards not-(L) being false. Thus, if Anna's evidence and evidential standards are such that Anna is permitted to have a very high credence in (L), they should also be such that Anna is not permitted to also have a very low credence in (L). But if Choice is true, then Anna's evidence and evidential standards point towards (L) being certainly true even though they also point towards (L) being almost certainly false. Thus, even if rationality is somewhat permissive in general, the first challenge for Choice is to explain how it can be so extremely permissive in particular.[15]

Even if rationality turns out to be extremely permissive, though, the second worry for Choice is that it licenses agents to reason in problematic ways. In IMPLICIT BIAS, for instance, recall that Choice permits Sally have any of these three combinations of attitudes:

T1 $\boxed{b}(p) \wedge \boxed{b}(E \nRightarrow p)$

T2 $\boxed{b}(p) \wedge \neg \boxed{b}(E \nRightarrow p)$

T3 $\neg \boxed{b}(p) \wedge \boxed{b}(E \nRightarrow p)$

Thus, Choice permits Sally to believe that Carl is rude and that her evidence does not actually support this (i.e.,"take 1"). But from Sally's own perspective, it is reasonable to ask: why is the evidence pointing away from what she believes? That is, why should Sally continue to believe that Carl is

---

[15]This objection leverages the standard claim (see e.g. White 2009) that evidence cannot point towards p while also pointing towards not-p. It is also worth noting that there are far less extreme versions of Intrapersonal Permissivism that do not run into the problem we are raising here. For instance, Jackson (forthcoming) defends a moderate version of Intrapersonal Permissivism according to which one agent, with one body of evidence and one evidential standard can be permitted to believe p in a low stakes case and suspend judgment about p in a high stakes case. Thus, unlike Extreme Intrapersonal Permissivism, Jackson has an explanation for why belief and suspension of judgment can both be permissible even though one's evidence and evidential standards remain the same: the stakes have changed drastically. Because Extreme Intrapersonal Permissivism permits agents to have radically different attitudes even though everything is held fixed, explanations like this one are not available.

rude when, by her own lights, her evidence does not support this? From Sally's point of view, the best explanation seems to be that the evidence is misleading. After all, if it really was pointing towards the truth, then Sally should follow the evidence where it leads and not believe that Carl is rude. Thus, by permitting Sally to have this combination of attitudes, Choice also licenses Sally to infer that her evidence is misleading. But no theory of rationality should license an agent to bootstrap her way to the conclusion that her evidence is misleading in this way (Horowitz, 2014).

Similarly, Choice permits Sally to believe that Carl is rude while also refusing to believe that her evidence does not actually support this (i.e., "Take 2"). While this might not seem problematic initially, recall that Sally has taken many implicit bias tests which give her good reason to think that her evidence does not actually support the claim that Carl is rude. From Sally's own perspective, then, why ignore these test results? After all, if Sally should take these results into account, then she should think that her evidence does not actually suggest that Carl has been impolite. From Sally's own perspective, the best explanation here seems to be that her implicit biases did not impact her judgment about Carl's rudeness. Thus, by permitting Sally to have this combination of attitudes, Choice licenses Sally to reason her way to the conclusion that she was immune from the effects of her implicit biases while she was interacting with Carl. But no theory of rationality should license an agent to form beliefs about their own psychology in this way.

Thus, proponents of Choice face two challenges. They must explain how rationality can be so extremely permissive, and also that such a permissive view does not license agents to reason in problematic ways. Neither challenge seems easy to meet, and both are avoided by Indeterminacy, which suggests that (in a case with the right structure) neither of these attitudes is determinately permissible, but at best indeterminately so.

5 Objections to priority

According to Priority, whenever the rules of rationality conflict, some consistent subset, identified on the basis of priority relations, must be privileged over the others. But how and why do the rules get ranked in a particular way? In IMPLICIT BIAS, for instance, what sorts of facts are such that Evidential Support gets assigned more (or less) importance than Inter-Level Coherence? Proponents of Priority must explain how conflicting epistemic rules get ranked

such that an agent is always permitted and required to satisfy one rule in particular.

Knoks (2020b) suggests that cases like IMPLICIT BIAS are under-described and that once a complete description is given, we can rely on the context to determine how much weight to give to each rule. For instance, once we know more about (i) Sammy's evidence for thinking that Carl is rude (e.g., how often they have interacted and what have those interactions have been like) (ii) how much evidence Sammy has regarding her own implicit biases and (iii) how much Sammy cares about acquiring true beliefs and avoiding error, then we will have the contextual information needed to know exactly how Evidential Support stacks up against Inter-Level Coherence.

A problem with this response is that, while contextual information might reveal how certain rules get ranked in some cases, context alone won't establish clear rankings in every case. Imagine, for instance, that Sammy has interacted with Carl three times. During the first interaction, it seemed to Sammy that Carl did not sincerely thank her for holding the door. During the second interaction, it seemed to Sammy that Carl didn't greet her after she waived hello. During the third interaction, Sammy noticed that Carl did not RSVP to her party invitation. Moreover, because she has taken a wide variety of implicit bias tests, Sammy knows that she is likely to misjudge Carl's behavior as being rude 41 percent of the time. Finally, Sammy values acquiring true beliefs and avoiding error to the same degree. Given this contextual information, should Evidential Support be given more weight than Inter-Level Coherence such that Sammy should believe that Carl is rude? Or should it be given less weight such that Sammy should not believe that Carl is rude? To our mind, the context alone doesn't provide definitive insight.

Things seem even less clear in LOGIC PROBLEM. Here it is stipulated that Anna is justified in thinking that there is exactly a 20 percent chance that she has arrived at the right answer about (L). It is also stipulated that (L) is a tautology. And let us suppose that Anna cares as much about having accurate credences as she does about avoiding inaccurate ones. Even so, this contextual information alone does not seems to shed much light on whether Probabilism should get more weight than Calibrationism (in which case Anna should be certain that (L) is true), or whether it should be the other way around (in which case Anna's credence in (L) should be .2).

The challenge of establishing clear priority rankings will also be difficult to meet, and it too is avoided by Indeterminacy, which says that in these cases,

there is no determinate fact of the matter about how the priority rankings get established.

6   Comparing two types of indeterminist theories

Indeterminacy provides a unified way of treating epistemic conflicts while also avoiding the worries discussed above. This is a good reason to prefer it to its rivals.

However, one important consequence of this paper is that there are many different ways of spelling out the Indeterminacy position. Knoks's work shows that it is important to attend to the distinction between rule schemas (such as the general support principle ES) and rule instances (such as the ESF and ESH). Similarly, in developing an indeterminacy-friendly position, we face a choice point concerning whether to think of resolutions of the indeterminacy in terms of schemas or instances. Leonard (2020) is naturally read as thinking of indeterminacy as between two norm-schemas. Our present analysis of the Indeterminacy position treats the indeterminacy as relating specific instances. Thus we may have a context whose ordering is $ESH \prec IRI \prec ESF$ where the two instances of the Evidential Support rule are separated in terms of their priority. In such a context, the resolution that generates the salient requirements is $\{ESF, IRI\}$. This is not an expressible resolution within a rule-schemas framework.

This is not to say that the approach that centers on rule schemas is obviously incorrect. There might be conceptual reasons to think that priority relations should hold between *general* rules and not their specific instances. Some may even feel that there is something incoherent-seeming about saying that the evidential support requirements have higher priority when they are applied to higher order claims about support relations, than they do when applied about specific propositions.

Be all that as it may, the approach in which priority relations hold between instances of rules is strictly more expressive. For this reason, we find it preferable.

In addition to this conceptual point, the Indeterminacy framework we built on the foundation of Knoks's system has modeling potential that is simply absent from a plain supervaluationist treatment of Indeterminacy. For one thing, our system allows for priority relations over rules, which leaves open the possibility that cases of conflicts can be resolved by some rules taking priority over others.

Of course, the flat-footed supervaluationist could reach out for a priority relation among epistemic rules, and then claim that one ought to supervaluate only over those rules that are not defeated. This move just incorporates within the supervaluationist story some key elements of our framework.

But while this maneuver would get them closer to our theory, there are still important outstanding differences. Consider this example:

> The Brando family follows the advice of three authorities: an academic expert, a priest for their religion, and their fitness coach. The Brandos view all of these authorities as reliable, but they view the academic and the priest as most reliable, and the tennis coach as slightly less reliable.

> Imagine now that the academic expert and the religious expert disagree on some issue: the academic expert says that masks help avoid certain infections; the priest says that they do not; the fitness coach says that purple masks help.

To model this situation, let r_a, r_p and r_f (respectively for "academic", "priest" and "fitness") be three rules corresponding to the three sources of authority. We suppose that r_a and r_p have priority over r_f but are incomparable with respect to each other. In this case, a supervaluationist account enriched with priorities would reason as follows: r_f is defeated by r_p, so the undefeated rules are the members of the set $\{r_a, r_p\}$. These are conflicting, and so we have to consider the two resolutions $\{r_a\}$ and $\{r_p\}$. The believer faces a tension between the claim that masks help and the conflicting claim that they do not.

This is progress, but it gets an important modeling detail wrong. Ideally, we would like the resolution that agrees with the academic to also be bolstered by the fitness coach's opinion. After all, this opinion was defeated by the priest, but when we think from the perspective of the academic expert, the priest's opinion is defeated. In other words, instead of a resolutions with r_a as its only member, we would like there to be a resolution corresponding to $\{r_a, r_f\}$.

Our system allows us to retrieve this distinction because of our distinction between proto-resolutions and resolutions. In the example, our proto-resolutions are $\{r_f\}$ and $\{r_p\}$. But as we move from proto-resolutions to resolutions, $\{r_f\}$ gets augmented with r_a. As a result, we end up with $\{r_a, r_f\}$ and $\{r_p\}$ as resolutions, just as we had hoped.

Our mechanism for turning proto-resolutions into resolutions is not perfect, and will itself need refinement in the face of even more complex problems.

Our point here is not that we have landed on the ultimate implementation of the idneterminacy position. Instead our point is that the default framework has the appropriate flexibility and expressive potential to capture a sophisticated version of the view that there is normative indeterminacy in the epistemic domain.

Bibliography

Brouwer, Thomas NPA (2014), 'A paradox of rejection', *Synthese*, 191(18), 4451–4464.

Caie, Michael (2012), 'Belief and indeterminacy', *Philosophical Review*, 121(1), 1–54.

Caie, Michael (2013), 'Rational probabilistic incoherence', *Philosophical Review*, 122(4), 527–575.

Christensen, David (1996), 'Dutch-book arguments depragmatized: Epistemic consistency for partial believers', *The Journal of Philosophy*, 93(9), 450–479.

Christensen, David (2007), '1. does murphy's law apply in epistemology? self-doubt and rational ideals', *Oxford studies in epistemology*, 1.

Christensen, David (2010), 'Higher-order evidence 1', *Philosophy and Phenomenological Research*, 81(1), 185–215.

Christensen, David (2016), 'Disagreement, drugs, etc.: From accuracy to akrasia', *Episteme*, 13(4), 397–422.

Conee, Earl (1987), 'Evident, but rationally unacceptable', *Australasian Journal of Philosophy*, 65(3), 316–326.

De Finetti, Bruno (1992), 'Foresight: Its logical laws, its subjective sources', in 'Breakthroughs in statistics', 134–174, Springer.

Egan, Andy and Elga, Adam (2005), 'I can't believe i'm stupid', *Philosophical Perspectives*, 19, 77–93.

Horowitz, Sophie (2014), 'Epistemic akrasia', *Noûs*, 48(4), 718–744.

Horty, John (2003), 'Reasoning with moral conflicts', *Noûs*, 37, 557–605.

Horty, John F. (2012), *Reasons as Defaults*, Oxford University Press.

Hughes, Nick (2019), 'Dilemmic epistemology', *Synthese*, 196(10), 4059–4090.

Jackson, Elizabeth (forthcoming), 'A defense of intrapersonal belief permissivism', .

Joyce, James M (1998), 'A nonpragmatic vindication of probabilism', *Philosophy of science*, 65(4), 575–603.

Knoks, Aleks (2020a), *Defeasibility in Epistemology*, Ph.D. thesis, University of Maryland.

Knoks, Aleks (2020b), 'Misleading higher-order evidence, conflicting ideals, and defeasible logic', *Ergo*, online.

Lasonen-Aarnio, Maria (2014), 'Higher-order evidence and the limits of defeat', *Philosophy and Phenomenological Research*, 88(2), 314–345.

Lasonen-Aarnio, Maria (2020), 'Enkrasia or evidentialism? learning to love mismatch', *Philosophical Studies*, 177(3), 597–632.

Leonard, Nick (2020), 'Epistemic dilemmas and rational indeterminacy', *Philosophical Studies*, 177(3), 573–596.

Littlejohn, Clayton (2018), 'Stop making sense? on a puzzle about rationality', *Philosophy and Phenomenological Research*, 96(2), 257–272.

Pettigrew, Richard (2016), *Accuracy and the Laws of Credence*, Oxford University Press.

Priest, Graham (2002), 'Rational dilemmas', *Analysis*, 62(1), 11–16.

Ross, Jacob (2010), 'Sleeping beauty, countable additivity, and rational dilemmas', *Philosophical Review*, 119(4), 411–447.

Roush, Sherrilyn (2009), 'Second guessing: A self-help manual', *Episteme*, 6(3), 251–268.

Schechter, Joshua (2013), 'Rational self-doubt and the failure of closure', *Philosophical Studies*, 163(2), 429–452.

Skyrms, Brian (1966), 'Choice and chance: An introduction to inductive logic', .

Sliwa, Paulina and Horowitz, Sophie (2015), 'Respecting all the evidence', *Philosophical Studies*, 172(11), 2835–2858.

Sorensen, Roy A (1987), 'Anti-expertise, instability, and rational choice', *Australasian Journal of Philosophy*, 65(3), 301–315.

Stapleford, Scott and McCain, Kevin (2020), 'Bound by the evidence', *Epistemic Duties: New Arguments, New Angles*.

Turri, John (2012), 'A puzzle about withholding', *The Philosophical Quarterly*, 62(247), 355–364.

White, Roger (2009), 'On treating oneself and others as thermometers', *Episteme*, 6(3), 233–250.

Worsnip, Alex (2018), 'The conflict of evidence and coherence', *Philosophy and Phenomenological Research*, 96(1), 3–44.