



# Rational requirements for suspended judgment

Luis Rosa<sup>1</sup> 

Published online: 2 March 2020  
© The Author(s) 2020

**Abstract** How does rationality bind the agnostic, that is, the one who suspends judgment about whether a given proposition is true? In this paper I explore two alternative ways of establishing what the rational requirements of agnosticism are: the Lockean–Bayesian framework and the doxastic logic framework. Each of these proposals faces strong objections. Fortunately, however, there is a rich kernel of requirements of agnosticism that are vindicated by both of them. One can then endorse the requirements that belong to that kernel without thereby committing oneself to the problematic implications that stem from either of the aforementioned proposals.

**Keywords** Rational requirements · Suspended judgment · Coherence · Probabilism

## 1 How does rationality bind the agnostic?

Theists believe that *God exists*, and atheists believe that *God does not exist*. Agnostics suspend judgment on the issue. There are many other things one might suspend judgment about: whether determinism is true, whether the world population this year exceeds the world population last year, whether one has been secretly followed, etc. But what are the rational requirements for suspended judgment? How does rationality bind the agnostic? In particular, does suspending judgment about whether a given proposition is true rationally require one to suspend judgment about other propositions as well?

The requirements that I have in mind here are requirements of *structural* rationality. They are about how our attitudes fit together in a coherent manner

---

✉ Luis Rosa  
fsopho@gmail.com

<sup>1</sup> CONCEPT, University of Cologne, Albertus-Magnus-Platz 1, 50923 Cologne, Germany

(Broome 1999; Christensen 2004). Throughout the paper, I will focus on *wide-scope* requirements of rationality, meaning that their main operator is the requirement operator, which ranges over entire combinations of doxastic attitude-ascriptions.<sup>1</sup> I will only assume that these are requirements in the *evaluative* sense: by failing to abide by these requirements, one becomes incoherent, and incoherence makes for irrationality. So one is only required to be/not to be a certain way *on pain of irrationality*.

Paradigmatic examples of requirements of this kind include those featuring ascriptions of *belief*. For example, where  $p$  logically entails  $q$  (or add ‘obviously so’ if you prefer):

- (B1) One is rationally required to be such that: if one believes that  $p$ , then one believes that  $q$ .
- (B2) One is rationally required to be such that: if one believes that  $p$ , then one does not believe that  $\neg p$ .<sup>2</sup>

That (B1) and (B2) are *wide-scope* requirements means, among other things, that there is more than one way to satisfy them. One can satisfy (B1) not only by believing that  $q$ , but also by refraining from believing that  $p$ . Either way, one would end up in a coherent doxastic state, as far as (B1) itself is concerned. (B1) and (B2) are both grounded on logical relations among propositions. (B1) stems from the fact that  $p$  logically entails  $q$ , and (B2) stems from the fact that it is logically impossible for both  $p$  and  $\neg p$  to be true.

According to *probabilism*, our credences must abide by the principles of the probability calculus (Christensen 2004; Pettigrew 2016). Where  $p$  logically entails  $q$ , any probability function  $Pr$  will be such that  $Pr(p) \leq Pr(q)$ .<sup>3</sup> Furthermore,  $Pr(\neg p) = 1 - Pr(p)$ , for any such function  $Pr$ . Accordingly, probabilism delivers the following wide-scope requirements for credal states (where  $p$  logically entails that  $q$ ):

- (P1) One is rationally required to be such that: one’s credence that  $q$  is the case is at least as high as one’s credence that  $p$  is the case.
- (P2) One is rationally required to be such that: if one has a credence of  $x$  that  $p$  is the case, then one has a credence of  $1 - x$  that  $\neg p$  is the case.

<sup>1</sup> See Broome (2007) for defense of wide-scope requirements, and Kolodny (2007) for criticism. Brunero (2010) and Way (2011) give what I take to be good responses to some of the objections/worries about wide-scope requirements. The points that I make in this paper do not really depend on whether coherence requirements are better interpreted as wide- or as narrow-scope requirements (the requirements of agnosticism that I am going to flesh out here can easily be reformulated as narrow-scope requirements).

<sup>2</sup> The conditionals that are under the scope of these requirement operators are to be read as a *material* conditionals. So, for example, (B2) is equivalent to the claim that one is rationally required not to believe that  $p$  and believe that  $\neg p$  at the same time. See MacFarlane (2004) and Steinberger (2016) for more comprehensive discussions about the connection between logical entailment and wide-scope norms or requirements that feature material conditionals.

<sup>3</sup> By ‘logically entails’ here I just mean the entailment relation from classical propositional logic. More generally, however, we can say the following: where  $p$  entails  $q$ ,  $Pr(p) \leq Pr(q)$  for any probability function  $Pr$  whose field or algebra encodes that entailment relation (in terms of subset relations).

To say that requirements (B1), (B2), (P1), (P2) are ‘grounded on’ (or ‘derived from’) logic and probability theory is not to say that logic and probability theory *by themselves* provide justification for them. The justification might have to do with considerations about the accuracy of doxastic states,<sup>4</sup> or considerations about the defeasibility of warrant or justification,<sup>5</sup> or with what is knowable a priori to rational subjects.<sup>6</sup> Either way, logical entailment and facts about probabilities will be a constitutive part of what grounds these requirements of rationality.

How could we manage to derive similar requirements of structural rationality for *agnostic* states from logic or probability theory? In this paper, I will explore two alternative ways of doing so. The first one—the ‘Lockean–Bayesian’ way—takes suspended judgment to be rationally tied to *middling credence*. Since the Lockean–Bayesian also takes the probability calculus to be normative for credences, wide-scope requirements for agnostic attitudes can then be derived from the principles of probability. The Lockean–Bayesian proposal will soon become controversial, however. The problem stems from Friedman’s (2013b) contention that some evidential states permit joint agnostic attitudes that do not correspond to permitted combinations of middling credences. In laying out the argument against the Lockean–Bayesian proposal, I will rely on the assumption that every violation of a coherence requirement such as the ones presented above involves failure to respond to one’s evidence in some way or other.

There is, however, an alternative proposal as to how to ground rational requirements for suspended judgment on logic: only it is *doxastic logic* (as opposed to the probability calculus). This proposal will allow us to theorize about the normativity of agnostic states by focusing on the following aspect of suspended judgment: that when one suspends judgment about whether  $p$  is true, both  $p$  and  $\neg p$  might turn out to be the case from one’s perspective. In contrast to the Lockean–Bayesian framework, the doxastic logic framework says that *it is* permissible for one to jointly hold the agnostic attitudes that Friedman (2013b) takes to be jointly permissible. But since the doxastic logic framework will bring consistency and closure requirements for belief in its wake, it will also be controversial.

I will then show that, despite this impasse, there is a rich kernel of rational requirements for agnostic states that both the Lockean–Bayesian and the doxastic logic frameworks agree on. These two frameworks are leading ways of representing the structure of rational doxastic attitudes in a formal setting. If there is a standoff between them (denying closure and consistency requirements is as bad as denying the permissibility of Friedman’s agnostic states), then, we can at least use the requirements that belong to that kernel to assess the coherence of agnostic states. Regardless of which of these frameworks turns out to be the best one at the end of the day, we will have a solid foundation for making such assessments and theorizing about the rationality of suspended judgment in a more theory-neutral fashion.

<sup>4</sup> See Joyce (1998), Easwaran and Fitelson (2015) and Pettigrew (2016).

<sup>5</sup> See Pollock (1986, 1990, ch. 9).

<sup>6</sup> See Smithies (2015) and Titelbaum (2015).

## 2 The Lockean–Bayesian framework

According to the Bayesian epistemologist, a coherent credal state must be properly represented by a (at least one) probability function. A subject who has a credence of  $x$  that  $p$  is the case and a credence of  $y > x$  that  $(p \wedge q)$  is the case thereby fails to abide by the requirements of rationality: there is no probability function that makes a conjunction more probable than one of its conjuncts.

So far, this is just the *Bayesian* speaking. The *Lockean–Bayesian* enters the scene when an attempt is made not only to ground rational requirements for credences on the probability calculus, but also requirements for the categorical attitudes of *belief*, *disbelief* and *suspended judgment*.<sup>7</sup> The rational connection between credences and categorical attitudes is then established through the following requirements:

- (L<sub>B</sub>) One is rationally required to be such that: one believes that  $p$  if and only if one has a high credence that  $p$  is the case;
- (L<sub>S</sub>) One is rationally required to be such that: one suspends judgment about  $p$  if and only if one has a middling credence that  $p$  is the case.<sup>8</sup>

That one has high credence that  $p$  is the case can be represented thus:  $Cr(p) > t$ .  $Cr$  is the credence function that represents one's credal state—it is a mapping from the propositions that one is opinionated about to values in the interval  $[0, 1]$ .  $t$  is a threshold (minimally  $t > 0.5$ ) above which credences start counting as belief-worthy.<sup>9</sup> Similarly, that one has a 'middling' credence that  $p$  is the case can be represented as follows:  $Cr(p) = m \in [1 - t, t]$ . The middling credence is the one that lies between  $t$ , that is, the threshold above which credences are normatively correlated with believing that  $p$ , and  $1 - t$ , that is, the threshold above which credences are normatively correlated with disbelieving that  $p$ /believing that  $\neg p$ .

The combination of probabilism, (L<sub>B</sub>) and (L<sub>S</sub>) gives rise to several rational requirements for suspended judgment, including the following (see the "[Appendix](#)" for details):

- (SS) One is rationally required to be such that: one suspends judgment about  $p$  if and only if one suspends judgment about  $\neg p$ .
- (SB) One is rationally required to be such that: if one suspends judgment about  $p$  then one does not believe that  $p$ .
- (SE) One is rationally required to be such that: if one suspends judgment about  $p$  and one has maximal credence that  $(p \equiv q)$ , then one suspends judgment about  $q$ .<sup>10</sup>

<sup>7</sup> Different versions of the so-called 'Lockean Thesis' are discussed in Foley (1993, ch. 4), Leitgeb (2014), and Pettigrew (2016, ch. 4).

<sup>8</sup> (L<sub>B</sub>) is not to be confused with the metaphysical claim that belief *just is* high credence—see Christensen (2004, pp. 84–85). Similarly for (L<sub>S</sub>).

<sup>9</sup> When interpreted in this way—that is, through  $Cr(p) > t$  instead of  $Cr(p) \geq t$ —(L<sub>B</sub>) does not include the view that one is required to believe that  $p$  iff one has *maximal* credence that  $p$ . Defenders of the Lockean Thesis usually take the threshold for belief to be less than 1—see Kyburg (1970), Foley (1993) and Hawthorne and Bovens (1999). I will briefly discuss the 'credence one' view in Sect. 7.

<sup>10</sup> '≡' is the symbol for material equivalence.

Suspended judgment is thought of as an attitude of *neutrality*: when one suspends judgment about  $p$ , one has a neutral stance regarding whether  $p$  or rather  $\neg p$  is the case (Sturgeon 2010; Friedman 2013a). Accordingly, (SS) says that you should not be neutral regarding whether  $p$  is true without also being neutral regarding whether  $\neg p$  is true, and vice-versa. (SB) just tells us that it is irrational to believe a proposition that one is agnostic about (let alone whether this is even possible). The maximal credence mentioned in (SE) is credence 1. So (SE) says that one should be agnostic about  $q$  if one is agnostic about  $p$  and is maximally confident that  $p$  and  $q$  are equivalent.

I take it that all these requirements are correct, and that they are very fundamental ones at that. So the Lockean–Bayesian has at least a *prima facie* plausible proposal regarding what the rational requirements for suspended judgment are. The following sections (3 and 4) will call the plausibility of this proposal into question, however.

### 3 Friedman’s permissibility claim

Friedman (2013b) claims that there are total evidential states that permit one to suspend judgment about a long conjunction of mutually independent, deeply contingent propositions, while at the same time suspending judgment about each of its conjuncts.<sup>11</sup> That is, there are total bodies of evidence that permit one to be in the following kind of doxastic state, where the  $p_i$  satisfy the conditions just mentioned, and ‘ $Sp_i$ ’ means that the subject suspends judgment about whether  $p_i$  is the case:

$$(F) \quad Sp_1 \wedge \dots \wedge Sp_n \wedge S(p_1 \wedge \dots \wedge p_n),$$

In making this claim, Friedman relies on the *absence of evidence* norm: ‘in the absence of evidence for or against an ordinary contingent proposition  $p$ , it is epistemically permissible to suspend judgment about  $p$ ’ (Friedman 2013b, p. 60).

Friedman presents the following case to illustrate the phenomenon. Ina is told that snowflakes will be collected from different parts of the globe.<sup>12</sup> Her task is to decide, for each selected snowflake  $x_i$ , whether  $x_i$  matches a certain shape (the shape is depicted in a drawing, or a digital image). Even though Ina has an average understanding of what snowflakes are, she does not have any relevant evidence to go about deciding what the shape of a snowflake is: ‘she has no idea what sorts of shapes they can take, how many different shapes they come in, the frequency with which a given shape occurs, or occurs with other shapes, and so on’ (Friedman 2013b, p. 62).

Given Ina’s evidential state, then, it would seem that she is at the very least *permitted* to suspend judgment about whether  $p_1$  is true, about whether  $p_2$  is true,

<sup>11</sup> The relevant notion of independence here is the notion of *probabilistic* independence:  $Pr(p_i | p_j) = Pr(p_i)$  for any  $j \neq i$  with  $i, j \in \{1, \dots, n\}$ . That a proposition is ‘deeply’ contingent means that ‘subjects will have no semantic guarantee of their truth’ and that ‘they simply describe mundane (possible) facts about the physical world’ (Friedman 2013b, p. 60).

<sup>12</sup> Friedman does not use the name ‘Ina’ herself—she uses the more neutral ‘S’ instead. As the reader will soon realize, however, there are already too many ‘S’s in my notation.

and so on, where  $p_i$  is the proposition that *snowflake  $x_i$  is a match*. And it also seems that she is permitted to suspend judgment about whether  $(p_1 \wedge p_2 \wedge p_3)$  is true. The size of the conjunction doesn't really matter here—Ina is just as well permitted to suspend judgment about any other arbitrarily large conjunction of propositions of the form  $p_i$  while suspending judgment about its conjuncts. These permissibility claims make sense. For all Ina knows, it might be that all snowflakes exhibit the relevant shape—but it might also be that their shapes vary wildly (she does not know it either way).

Friedman's target is not exactly the Lockean–Bayesian view on suspended judgment that I introduced above, again:

(L<sub>S</sub>) One is rationally required to be such that: one suspends judgment about  $p$  if and only if one has a middling credence that  $p$  is the case.

Rather, she is aiming at a *reductionist* thesis according to which suspended judgment about  $p$  is just a matter of having a middling credence towards  $p$ . But her objection is obviously relevant to the current investigation. For one of the alleged differences between suspended judgment and middling credence here stems from the fact that, while there are evidential situations that permit one to be in a state of type (F), no corresponding state featuring middling credences can be represented by a single probability function. For we are assuming, remember, that there is probabilistic independence between the  $p_i$  that feature in the target state of type (F). The probability of  $(p_1 \wedge \dots \wedge p_n)$ , then, is just the product of the individual probabilities of the  $p_i$ . But if a credal state features middling credences in each of  $p_1, \dots, p_n$ , and  $(p_1 \wedge \dots \wedge p_n)$  respectively—and the  $n$  is large enough to make  $t^n$  fall below the middling credence interval  $[1 - t, t]$ —then *there isn't* a probability function that represents that credal state.<sup>13</sup> Contrapositively, if *there is* a probability function that represents that credal state, then not all of its credences are middling (we just need to make  $n$  large enough).

The next section shows more precisely how Friedman's case against the reductionist thesis also makes trouble for the Lockean–Bayesian view that I presented in Sect. 2.

#### 4 From evidential permissibility to coherence

Friedman's claim is about the *evidential* permissibility of states of type (F), or permissibility relative to one's total evidence. It belongs therefore to the realm of 'substantial' rationality. How is the transition to be made from there to the requirements of 'structural' rationality, that is, coherence requirements?<sup>14</sup> If there

<sup>13</sup> Why  $t^n$ ? Think of it like this:  $t$  is the *highest* credence that still correlates with suspended judgment—so that is the best chance that a large conjunction has to have its probability still in the interval  $[1 - t, t]$ .

<sup>14</sup> This way of asking the question seems to imply that the requirement that one proportions one's attitudes to one's evidence isn't itself a coherence requirement. I am aware that someone could call this into question. But I am using the terms 'substantial' and 'structural rationality' just to flag an *apparent* contrast between two kinds of requirements. I won't take a stand on whether these appearances are misleading here.

weren't such a connection, the following move would be open to the Lockean–Bayesian: 'Maybe Ina's *evidence* permits her to suspend judgment about all of  $p_1, \dots, p_n$  and  $(p_1 \wedge \dots \wedge p_n)$  at the same time—but it is still *incoherent* for her to do that. In order for her to achieve coherence, she must revise at least some of her attitudes'.<sup>15</sup>

The following thesis would make the target transition fairly straightforward, and it would thereby block this move:

- (1) Every violation of a requirement of structural rationality involves violation of *some* requirement of substantial rationality.

If Friedman (2013b) is right, then substantial rationality permits Ina to be in the doxastic state  $(Sp_1 \wedge \dots \wedge Sp_n \wedge S(p_1 \wedge \dots \wedge p_n))$ , even when  $n$  is very large (say, because she was asked to give her opinion about 50 randomly selected snowflakes). That is, Ina does not violate any requirement of substantial rationality when she is in that state. Given (1), it follows that Ina does not violate any requirement of structural rationality when she is in that state (by universal instantiation and *modus tollens*). According to the Lockean–Bayesian, however, Ina's doxastic state *does* violate a coherence requirement: either it violates  $(L_S)$  or, if it doesn't, there is no probability function that represents that state (assuming, again, that the  $p_i$  are deeply contingent and mutually independent, and that  $n$  is large enough). Either way, Ina's state would count as incoherent by the Lockean–Bayesian's standards.

Friedman (2013b) has made a good case for the claim that a subject in a situation such as Ina's is permitted to simultaneously suspend judgment about each of  $p_1, \dots, p_n$ , and  $(p_1 \wedge \dots \wedge p_n)$ . (1) is furthermore a plausible view: it would be surprising if one could hold incoherent beliefs and yet be normatively in the clear as per one's total evidence or reasons. In order to endorse (1), we need not commit ourselves to the view that the requirements of structural rationality simply reduce to requirements of substantial rationality, or that the former ones can be dispensed with given that we have the latter ones (see Kolodny 2007; Kiesewetter 2017, Ch. 6–7). We could accept (1) and still take coherence talk to be an invaluable heuristics or means of making assessments about substantial rationality (lack of coherence signals lack of responsiveness to evidence or reasons—we might not know more exactly how the subject fails to be responsive to it, but we know that she fails to do so *somehow*). And we could also accept (1) while holding that rationality is a function of two different things: coherence and responsiveness to the evidence. Both of these views are live options, and I won't try to decide between them here. The point is just that (1) seems to be true, and not just trivially so (it seems that *there are*

<sup>15</sup> The Lockean–Bayesian would thereby embrace the existence of *rational dilemmas* where no funny higher-order phenomena is involved. Easwaran and Fitelson (2015) would reject this kind of rational dilemma: if all the attitudes of a doxastic state are properly responsive to the evidence, then there is a probability function that represents that doxastic state (and it is thereby coherent). I have the impression that many epistemologists will side with Easwaran and Fitelson here, at least when it comes to doxastic states whose attitudes are all of the same order.

true coherence requirements, whatever their truth ultimately depends upon).<sup>16</sup>

Worsnip (2018) rejects (1) when it is interpreted with full generality. He thinks that the same body of evidence can sometimes support an attitude  $Dp$  while at the same time supporting the belief that the subject's evidence does not support  $Dp$ . So the subject's evidence simultaneously supports attitudes that are not mutually coherent (Worsnip calls this *inter-level* incoherence). This phenomenon purportedly happens when the subject possesses misleading higher-order evidence. As Worsnip admits (2018, Section 4), however, it is hard to give 'knockdown' intuitive examples where the subject's total evidence simultaneously supports (inter-level) incoherent attitudes. Addressing Worsnip's contention would take me too far away from this paper's goals. For my present purposes, it will be enough to restrict (1) to coherence requirements that feature doxastic attitude ascriptions that are on the same level of description. In particular, one can restrict it to those requirements that are grounded on the formal-logical connections between the contents of the doxastic attitudes involved, i.e., requirements such as (B1), (B2), (P1) and (P2) from above. (It is possible for both  $p$  and *the subject's evidence does not give support to  $p$*  to be true at the same time, and the same probability function can assign high probability to both. By contrast,  $p$  and  $\neg p$  cannot both be true, and there is no probability function that assigns high probability to both.)

Notice, furthermore, that those who think that it can be rational for one to hold inconsistent beliefs (see Christensen 2004) are not thereby committed to denying (1). It might appear as if their diagnosis of Preface-like situations does commit them to the denial of (1): every single belief the subject has is supported by the evidence, and yet the subject's beliefs are mutually inconsistent. But what these philosophers are effectively saying is that we are *not* rationally required to be logically consistent. So there would be no violation of a requirement of structural rationality in Preface-like situations, which would make (1) trivially true rather than falsify it (the antecedent that is under the scope of its quantifier would not be satisfied in those cases).<sup>17</sup>

This argument from (1) and Friedman's permissibility claim against the Lockean–Bayesian account is plausible enough to warrant searching for another account of the structural requirements of agnosticism. More familiar objections to (different versions of) the Lockean Thesis add up to this case against the Lockean–Bayesian framework.<sup>18</sup> In the next two sections, I will explore an alternative

<sup>16</sup> Kiesewetter agrees with (1), but he doubts that there are any structural requirements of rationality to begin with. He argues that 'both narrow- and wide-scope interpretations of structural requirements face significant problems quite independently of whether these requirements are understood as being normative' (2017, p. 126). I do not think that the problems that different versions of these requirements face constitute good reason to deny that *there are any*, however.

<sup>17</sup> Also, if there are *no* requirements of rationality whatsoever, then there are no violations of rational requirements—which would again make (1) true, albeit trivially so. For more discussion on the general question whether incoherence always involves failure to respond to reasons (in both the theoretical and practical domains), see Way (2018). I thank an anonymous referee for these important observations.

<sup>18</sup> See, for example, Buchak (2014), who takes belief and high credence to be responsive to different features of one's evidence. See also Littlejohn (2015) and Staffel (2016).

framework that seeks to ground rational requirements on *doxastic logic*, as opposed to the probability calculus. That framework will vindicate the permissibility of states of type (F), even when arbitrarily large conjunctions are concerned.

## 5 Doxastic logic and rational requirements

Doxastic logics are commonly interpreted as logics of ideally rational belief, as opposed to belief *simpliciter*.<sup>19</sup> The standard semantics for doxastic formulas here is known as *possible-worlds semantics*. The belief-operator ( $B$ ) is treated here like any other necessity operator, in that the truth-conditions for ascriptions of belief are given by:

(Bel)  $M, w \models Bp$  iff for every  $v \in R(w)$ ,  $M, v \models p$ .

That is,  $Bp$  is true in a possible world  $w$  according to model  $M$  if and only if  $p$  is true in all the possible worlds that are doxastically accessible to the subject in  $w$  (the subject, that is, whose situation is represented by model  $M$ ). The model is a tuple  $M = \langle W, R, V \rangle$ , where  $W$  is a set of possible worlds,  $R$  is an accessibility relation on  $W$  and  $V$  is a valuation function that tells us which atomic formulas are true/false in which possible worlds.<sup>20</sup> That  $v \in R(w)$ , or that  $v$  is ‘doxastically accessible’ to the subject in  $w$ , just means that everything the subject believes in  $w$  is true in  $v$ . For all the subject believes in  $w$ , she might as well be in  $v$ .  $R(w)$  is the total set of worlds that are doxastically accessible to the subject in  $w$ —it sums up everything the subject believes in  $w$ .

$B$  represents ideally rational belief in the following sense:

- (a) If  $p_1, \dots, p_n$  logically entail  $q$  then, for any model  $M$  and possible world  $w$ , either  $(Bp_1 \wedge \dots \wedge Bp_n)$  is not true at  $w$  or, if it is, then so is  $Bq$  (if  $M, w \models Bp_1 \wedge \dots \wedge Bp_n$  then  $M, w \models Bq$ ).
- (b) Where  $p_1, \dots, p_n$  are jointly inconsistent, there is no model  $M$  and possible world  $w$  such that  $M, w \models Bp_1 \wedge \dots \wedge Bp_n$ .

<sup>19</sup> See Hintikka (1962) for a seminal work on this. As both Stalnaker (2006, p. 172) point out, however, there is more than one way of interpreting the theoretical constructs of doxastic logic: (a) they are supposed to model or represent the situation of an ideally rational believer who has unlimited or superhuman logical capacities, *versus* (b) they are supposed to model or represent the situation of human cognizers like you and me, but the logic is concerned with what is *implicit* in our total doxastic states, or simply what follows from our beliefs (see also Yap 2014). And so, for example, in order for  $q$  to be true in all the possible worlds that are not ruled out by what you believe (that is, in order for  $Bq$  to be true about your situation), you don’t have to have inferred that  $q$  from the beliefs that you actually have. It is just that an ideal reasoner would have done so (a), or that  $q$  is already implicit in all the things you believe (b). Notice that neither of these interpretations purports to say what it is for one to believe a proposition. I stick to the ideally-rational-belief interpretation because this makes the connection between doxastic logic and the requirements of rationality more straightforward/easier to make than the implicit-belief interpretation.

<sup>20</sup> An atomic formula  $p$  is true in  $w$  according to model  $M$ , or  $M, w \models p$ , when  $w \in V(p)$ .  $V$  is a mapping from formulas to sets of possible worlds. The truth-values of complex formulas are then computed using the usual compositional rules: (i)  $M, w \models \neg\phi$  iff it is not the case that  $M, w \models \phi$ , (ii)  $M, w \models (\phi \wedge \psi)$  iff  $M, w \models \phi$  and  $M, w \models \psi$ , etc.

These features of the  $B$  operator are supposed to capture the *coherence* dimension of ideally rational beliefs (more on (a) and (b) below).<sup>21</sup> Arguably, there is yet another dimension to ideal rationality: responsiveness to evidence. One could take the latter property into account by adding yet another accessibility relation to these models, which would then tell us what is the subject's total evidence in each situation. But I will leave these details about how to represent evidence responsiveness open for now.

The thesis that I am going to explore in this section and the next one is that a coherence requirement of the form:

- (CR) One is rationally required to be such that: if one satisfies all of  $\Phi_1, \dots, \Phi_n$  respectively, then one satisfies  $\Psi$

is true whenever  $\Phi_1, \dots, \Phi_n$  entail  $\Psi$  in doxastic logic (which is interpreted again as the logic of ideally rational belief).<sup>22</sup>

The entailment relation that I will deploy here relies only on the *seriality* of the doxastic accessibility relation  $R$ . The relevant system of doxastic logic will be at least strong as  $D$ —nothing stronger than that is needed for my present purposes.<sup>23</sup> So for every possible world  $w$  in  $W$  of a doxastic model  $M = \langle W, R, V \rangle$ , there is at least one possible world  $v$  in  $W$  such that  $v \in R(w)$ . This guarantees that there is no model  $M$  and possible world  $w$  such that  $M, w \models (Bp \wedge B\neg p)$ .

If this framework is all about ideally rational beliefs, how does it latch onto the doxastic lives of real thinkers like you and me? Roughly like this: take the propositions that you believe to be true in a given situation and isolate the possible worlds where those propositions are true. *Whatever else* is true in all of those worlds is also believed by your ideally rational counterpart.  $M, w \models Bp$  is not to be interpreted as saying that *you* believe that  $p$  in  $w$ , even though the relevant model is in a sense about you (we have used the propositions that you believe to be true in order to start building the model). Rather,  $M, w \models Bp$  is to be interpreted as saying

<sup>21</sup> The idea that there is a connection between the entailment relation from epistemic/doxastic logic and norms of rationality was already made by Hintikka himself (see especially 1962, §2.6–§2.10). In Hintikka's seminal work, the very notion of *consistency* in the logics of knowledge and belief has to do with defensibility of a subject's reporting what she does or does not believe/know. For example, that there is inconsistency between the formulas  $Bp, B(p \supset q)$  and  $\neg Bq$  of doxastic logic means (among other things) that a person whose cognitive state satisfies those three formulas will—if *she is reasonable or rational*—be persuaded to change that doxastic state by merely attending to certain entailment relations. She will then either come to believe that  $q$  or abandon one of her previous beliefs (either the belief that  $p$  or the belief that  $(p \supset q)$ ). If the subject is not subject to rational criticism in this manner, however, then the doxastic formulas that describe her cognitive state are not inconsistent.

<sup>22</sup>  $\Phi_1, \dots, \Phi_n$  and  $\Psi$  are used here as placeholders for *doxastic formulas* of a set  $DF$  that is recursively defined as follows (where  $PL$  is the language of propositional logic): (1) if  $\phi \in PL$  then  $B\phi \in DF$ , (2) if  $\Phi \in DF$ , then  $\neg\Phi \in DF$ , (3) if  $\Phi, \Psi \in DF$ , then  $(\Phi \wedge \Psi), (\Phi \vee \Psi), (\Phi \supset \Psi) \in DF$ . ' $\supset$ ' is the symbol for the material conditional. Notice that  $DF$  only contains doxastic formulas that do not have doxastic formulas for their contents. I want to reduce complexity at this point and leave the requirements of inter-level coherence for future investigation. This is in tandem with the observations I made in the previous section about (1).

<sup>23</sup> See Hughes and Cresswell (1968) for an overview of the different systems of modal logic, from K to S5.

that an ideally rational cognizer who has the beliefs that you have also believes that  $p$  in  $w$ .

In the next section I define a suspended judgment-operator using the language of standard doxastic logic and I explore the requirements on agnostic states that stem from this framework.

## 6 The suspension operator and the requirements of agnosticism

Let me now define a suspended judgment-operator  $S$  as follows:

(Def)  $Sp =_{def.} (\neg Bp \wedge \neg B\neg p)$ .

The belief-operator  $B$  has a dual  $C$ , meaning that  $Cp$  and  $\neg B\neg p$  are logically equivalent.  $C$  is the doxastic possibility-operator. The following is then a theorem of doxastic logic enriched with (Def):

(EQ)  $Sp \equiv (Cp \wedge C\neg p)$ .

That  $M, w \models Cp$  means that, as far as what the ideally rational subject believes in  $w$  goes, it is possible that  $p$  (the ‘ideally rational subject’ being the ideal counterpart of the subject whose situation we are theorizing about). That is: her beliefs do not rule out the possibility that  $p$  is the case. So  $M, w \models Sp$  means that, as far as what the ideally rational subject believes in  $w$  goes, it is possible that  $p$ , but it is also possible that  $\neg p$ . Relatedly, when I suspend judgment about, e.g., whether *Ashanti is in Zimbabwe*, I can express my agnostic stance toward that proposition by saying ‘Ashanti might be in Zimbabwe right now, but she might also not be’.<sup>24</sup>

The idea behind using (Def) in doxastic logic, then, is that this will allow us to theorize about the rationality of suspended judgment by thinking of situations where both a proposition and its negation are doxastically possible to an ideal counterpart of the subject. We are only interested in the rational requirements of agnosticism that will stem out of this framework. (Def) is not an analysis of what suspended judgment is, nor is it committed to any such particular analysis—just like the truth-conditions for the  $B$  operator do not constitute an analysis of the notion of belief. In particular, the idea here is not to reduce suspended judgment to absence of belief and disbelief toward a proposition. As already pointed out by many others, this is a very problematic view.<sup>25</sup> Nor could that be the idea, since  $M, w \models Bp$  does not mean that you believe that  $p$  in  $w$  either, as we saw in the previous section (even though, again, the model that verifies it is still in a sense about you).<sup>26</sup>

With the suspended judgment operator defined in this way, the following consequence relations will hold in any system of doxastic logic at least as strong as

<sup>24</sup> See Huemer (2011, p. 6) for a similar observation—even though Huemer uses ‘may’ instead of ‘might’.

<sup>25</sup> See, for example, Wedgwood (2002) and Friedman (2013a).

<sup>26</sup> As far as the current way of using (Def) goes, it might even turn out that suspended judgment is not a propositional attitude at all, as suggested by Friedman (2017).

$D$  (now I use ‘ $\models$ ’ for the *consequence relation* in  $D$ —see the “Appendix” for more details):

- (D1)  $Bp, Sq \models S(p \wedge q)$
- (D2)  $Sp, Sq, \neg B\neg(p \wedge q) \models S(p \wedge q)$
- (D3)  $S(p \wedge q), Bp \models Sq$
- (D4)  $B\neg p, Sq \models S(p \vee q)$
- (D5)  $Sp, Sq, \neg B(p \vee q) \models S(p \vee q)$
- (D6)  $S(p \vee q) \models Sp \vee Sq$
- (D7)  $B(p \supset q), Sq, \neg B\neg p \models Sp$
- (D8)  $Bp, Sq \models S(p \supset q)$
- (D9)  $S(p \supset q), S(p \supset \neg q) \models Sq$
- (D0)  $B(p \equiv q), Sp \models Sq$

To each of these entailment relations there correspond a rational requirement—the requirement that one be such that: if one satisfies all of the premises of that entailment relation, then one satisfies its conclusion. Consider (D1) as an example. The relevant requirement here is:

- (R1) One is rationally required to be such that: if one believes that  $p$  and one suspends judgment about  $q$ , then one suspends judgment about  $(p \wedge q)$ .

Suppose that I believe that *I am hungry now*, but I suspend judgment about whether *I will still be hungry two hours from now* (I’m not sure I will be able to grab a bite). In order for me not to violate (R1), then, my attitude towards *I am hungry now and I will still be hungry two hours from now* must be an equally agnostic one.

Some of the entailment relations above include what we might call ‘anti-exception clauses’ among their premises. As an example, consider (D2). The anti-exception clause here is  $\neg B\neg(p \wedge q)$ . We might initially hypothesize that there is a requirement that one suspends judgment about a conjunction whenever one suspends judgment about its conjuncts. But there are exceptions to this. Suppose that  $p$  and  $q$  are mutually exclusive within the subject’s space of doxastic possibilities: for all  $v \in R(w), M, v \models \neg(p \wedge q)$ . For example,  $p$  might be the proposition that *Lucy’s pet is a dog*, and  $q$  the proposition that *Lucy’s pet is a cat* (she only has one pet). It is not incoherent for one to suspend judgment about  $p$ , suspend judgment about  $q$  and *not* suspend judgment about  $(p \wedge q)$  here. The anti-exception clause in (D2), then, makes sure that this possible exception is properly taken into account. The requirement it gives rise to says that one is required to be such that: one suspends judgment about the conjunction if one suspends judgment about each of its conjuncts *and one does not reject their conjunction*. If one does reject the conjunction, however, one is already in a doxastic state that abides by that requirement.

In addition to the requirements corresponding to entailment relations D1–D0, the Doxastic Logic framework also gives us (SS) and (SB), just like the Lockean–Bayesian framework does:

- (SS) One is rationally required to be such that: one suspends judgment about  $p$  if and only if one suspends judgment about  $\neg p$ .

(SB) One is rationally required to be such that: if one suspends judgment about  $p$  then one does not believe that  $p$ .

For, given (Def),  $(Sp \equiv S\neg p)$  and  $(Sp \supset \neg Bp)$  are theorems of doxastic logic (albeit under a different mode of presentation).<sup>27</sup> In contrast to the Lockean–Bayesian framework, however, this proposal is compatible with the rational permissibility of states of type (F), even when they feature attitudes of suspended judgment toward arbitrarily large conjunctions. For, given only that  $p_1, \dots, p_n$  are mutually independent, deeply contingent propositions, it is not the case that  $Sp_1, \dots, Sp_n \models \neg S(p_1 \wedge \dots \wedge p_n)$ . There is after all a doxastic model that satisfies  $Sp_1 \wedge \dots \wedge Sp_n \wedge S(p_1 \wedge \dots \wedge p_n)$ . So the Doxastic Logic framework does not generate a requirement that one avoids states of type (F), which should please those who agree with Friedman (2013b).

Arguably, however, making room for Friedman’s permissibility claim comes at a price. The next section discusses some of the problems for the putative connection between doxastic logic and the requirements of rationality in this framework.

## 7 Closure and consistency requirements

The coherence requirements of agnosticism from the previous section are based on standard doxastic logic (any doxastic logic at least as strong as  $D$ ). That logic is interpreted as the logic of ideally rational belief. The idea behind this proposal is: if the ideally rational agent never ends up in a given doxastic state (there is no model that verifies the conjunction of certain doxastic formulas), then one must avoid that doxastic state in order to be coherent. For example, standard systems of doxastic logic tell us that  $B(p \wedge q)$  entails that  $\neg B\neg p$ , ergo that no model verifies the formula  $B(p \wedge q) \wedge B\neg p$ . Accordingly, it is incoherent for one to believe that  $(p \wedge q)$  and that  $\neg p$  at the same time. If one is in that state, then one should drop at least one of those beliefs (on pain of irrationality).

By deploying this method, however, one would seem to be committed to *all* the purported requirements that correspond to entailment relations/theorems of doxastic logic. And doxastic logic says that (a)  $B(T)$  is valid, where  $T$  is any tautology of propositional logic, and it also says that (b)  $Bp_1, \dots, Bp_n \models Bq$  when  $p_1, \dots, p_n$  entail that  $q$  in propositional logic. So one would be rationally required to believe all tautologies, and also required to be such that, if one believes each of  $p_1, \dots, p_n$  respectively, then one believes that  $q$  (assuming again that  $p_1, \dots, p_n$  logically entail  $q$ ). The Lockean–Bayesian framework also delivers a requirement to believe tautologies—but it *does not* deliver the requirement that one’s beliefs be closed under logical entailment in this way. Nor does it deliver the requirement that one’s beliefs be logically consistent. And one might think that the Lockean–Bayesian has an advantage over the doxastic logic proposal here, because it does not look as if

<sup>27</sup> The original modes of presentation being ‘ $((\neg Bp \wedge \neg B\neg p) \equiv (\neg B\neg p \wedge \neg B\neg\neg p))$ ’ and ‘ $((\neg Bp \wedge \neg B\neg p) \supset \neg Bp)$ ’ respectively.

one is rationally required to be consistent and to have one's beliefs closed under entailment.

Worries about consistency/closure requirements stem from intuitive assessments of situations such as the one depicted in the Preface Paradox (Makinson 1965): the subject believes each of a large series of propositions while believing that at least one of them is false. After all, the subject is aware that she is a fallible human being/that she eventually gets things wrong like everybody else. The intuitive assessment here would then be that the subject is rationally permitted to be in that kind of doxastic state, even though it is an inconsistent one.<sup>28</sup> So subjects would not be rationally required to be logically consistent, and their beliefs would not be required to be closed under logical entailment.

Of course, this assessment of the situation is not uncontroversial.<sup>29</sup> Various epistemological views are in conflict with it—including variants of the Lockean–Bayesian view (they differ from  $(L_B)$ ). For example, one might take belief to be *maximal* credence (credence one), perhaps in such a way that belief is sensitive to context. This view is also committed to the consistency and closure requirements for belief. There are different ways of motivating this view—see Greco (2015) for one such example.<sup>30</sup> Or consider Leitgeb's (2014) *stability theory*, which attempts to conciliate the closure/consistency requirements of belief with the Lockean thesis *without* taking the probabilistic threshold of belief to be the maximal one. According to this theory, belief is again sensitive to context, and the relevant threshold is determined by a stable probability of the subject's set of doxastically accessible scenarios ('stable' in the following sense: that set has probability  $> 0.5$  conditional on any subset of scenarios that is consistent with it).<sup>31</sup>

Finally, consider 'pluralists' or 'dualists', who think that belief and credence are both equally fundamental attitudes, meaning that neither is reducible to the other (see Jackson 2019). Pluralists seem to think that believing that  $p$ —in contrast to having high credence that  $p$  is the case—always involves ruling out the possibilities where  $p$  is false (Buchak 2014; Ross and Schroeder 2014).<sup>32</sup> And this is exactly what underlies the behavior of the  $B$  operator in doxastic logic.

<sup>28</sup> See Christensen (2004) for discussion, especially Chapters 3 and 4, which take a critical look on (among others) Pollock's (1986) defense of 'cogency'. See also Foley (1993) and Hawthorne and Bovens (1999).

<sup>29</sup> Some friends of consistency who disagree with it include Nelkin (2000), Ryan (1991) and Douven and Williamson (2006).

<sup>30</sup> See also Levi (1991), Clarke (2013), Greco (2015) and Dodd (2016). Greco's (2015) proposal is that credences are distributed over the set of possibilities that are doxastically accessible to one. Belief sets the tone for credence. For each state  $w$  that the subject might be in, then,  $Cr(R(w)) = 1$ . But  $R(w)$  is just the sum-up of what the subject believes, so belief is credence one. He presents an argument by analogy for this view and tries to defuse worries about the fact that we are not always maximally confident in what we believe to be true. His response is roughly that the possibilities that we treat as 'live' is sensitive to practical factors and other situational factors.

<sup>31</sup> For a critical take on Leitgeb's stability theory, see Staffel (2016).

<sup>32</sup> Here is Buchak (2014, p. 286): 'When an agent believes  $p$ , she in some sense rules out worlds in which *not-p* holds. The truth of *not-p* is incompatible with the attitude she holds towards  $p$  (though it is not incompatible with her holding that attitude, since she may be mistaken). On the other hand, having a

So, again, there are defensible epistemological views that do not conflict with (or they even vindicate) the requirements that stem from standard doxastic logic. But the fact that many epistemologists reject the consistency/closure requirements for belief prevents me at this point from just relying on doxastic logic as a guide to the rational requirements of agnosticism in the way sketched above. In the next section, then, I will focus on what the Lockean–Bayesian and the doxastic logic frameworks have in common, rather than on how they differ, and which one is the best framework overall (an issue I cannot hope to settle in this paper).

## 8 Requirements in common

I have pointed out in Sect. 6 that both the Lockean–Bayesian and the doxastic logic frameworks vindicate requirements (SS) and (SB). But this isn't a terribly comprehensive answer to this paper's initial questions: What are the rational requirements for doxastic states involving suspended judgment? In particular, does suspending judgment about whether a given proposition is true rationally require one to suspend judgment about other propositions as well?

Fortunately, however, there are other requirements that both the Lockean–Bayesian and the doxastic logic frameworks agree on. Consider again the entailment relation from doxastic logic:

$$(D2) \quad Sp, Sq, \neg B\neg(p \wedge q) \models S(p \wedge q).$$

According to the putative link between (the entailment relation of) doxastic logic and the requirements of rationality from Sects. 5 and 6, then:

$$(R2) \quad \text{One is rationally required to be such that: if one suspends judgment about } p, \text{ suspends judgment about } q, \text{ and does not believe that } \neg(p \wedge q), \text{ then one suspends judgment about } (p \wedge q).$$

It turns out that the Lockean–Bayesian framework will also endorse (R2). This is so due to the presence of the 'anti-exception' clause in the antecedent of the conditional in (R2): that one does not believe that  $\neg(p \wedge q)$ . Suppose that I suspend judgment about whether *it is going to rain tomorrow*, and also about whether *it is going to rain the day after tomorrow*. Suppose I also end up suspending judgment about whether *it is going to rain tomorrow and the day after tomorrow*. The Lockean–Bayesian will then say that I cannot go on like this. For each particular day  $d_i$  from now on, I might suspend judgment about whether it is going to rain that day, or  $Rain(d_i)$ —but I should not suspend judgment about whether *it is going to rain everyday from now on*. At some point, even though my credences toward  $Rain(d_n)$  and  $(Rain(d_1) \wedge \dots \wedge Rain(d_{n-1}))$  respectively are both middling, my credence toward  $(Rain(d_1) \wedge \dots \wedge Rain(d_n))$  must not be middling anymore (on pain of

---

Footnote 32 continued

particular credence in  $p$ , at least if it is not 0 or 1, does not rule out either the  $p$ -worlds or the *not- $p$*  worlds'.

incoherence). So suspending about  $(\text{Rain}(d_1) \wedge \dots \wedge \text{Rain}(d_n))$  fails to cohere with the previous middling, agnostic attitudes. But that is exactly the point where I am not satisfying the anti-exception clause anymore: once my credence in  $(\text{Rain}(d_1) \wedge \dots \wedge \text{Rain}(d_n))$  has gone lower than  $1 - t$ , the coherent thing for me to do is to believe that  $\neg(\text{Rain}(d_1) \wedge \dots \wedge \text{Rain}(d_n))$ . If I believe that  $\neg(\text{Rain}(d_1) \wedge \dots \wedge \text{Rain}(d_n))$ , however, I am already abiding by (R2)—in the sense that I am not violating it—since I do not satisfy all the conditions that are mentioned in the antecedent of the conditional that is embedded in it.

Similarly, consider the entailment:

$$(D7) \quad B(p \supset q), Sq, \neg B\neg p \models Sp.$$

Accordingly, the doxastic logic framework gives us the requirement:

- (R7) One is rationally required to be such that: if one believes that  $(p \supset q)$ , suspends judgment about  $q$  and does not believe that  $\neg p$ , then one suspends judgment about  $p$ .

Again, the Lockean–Bayesian framework will also endorse (R7), seeing as (R7) has the exception clause that one does not believe that  $\neg p$ . The Lockean–Bayesian will disagree, however, with requirements (R1), (R3), (R4), (R6) and (R8), which stem from (D1), (D3), (D4), (D6) and (D8) respectively. None of these requirements have the relevant anti-exception clauses. Consider (R3) as an example:

- (R3) One is rationally required to be such that: if one suspends judgment about  $(p \wedge q)$  and one believes that  $p$ , then one suspends judgment about  $q$ .

Since *there is* a probability function  $Pr$  such that  $Pr(p \wedge q) \in [1 - t, t]$ ,  $Pr(p) > t$  and  $Pr(q) > t$ , however, the Lockean–Bayesian will say that it is rationally permissible for one to suspend judgment about  $(p \wedge q)$  and believe that  $p$  while also believing that  $q$  (at least assuming again that the threshold required for belief isn't maximal).

A recipe for fleshing out other mutually accepted requirements, then, is to include the relevant anti-exception clauses whenever they are missing from the requirements that stem from the doxastic logic framework (they are 'missing' from the Lockean–Bayesian point of view). To continue with the example of (R3), the Lockean–Bayesian and the doxastic logic frameworks will agree on the following modification of it:

- (R3\*) One is rationally required to be such that: if one suspends judgment about  $(p \wedge q)$ , believes that  $p$  and does not believe that  $q$ , then one suspends judgment about  $q$ .

Similar anti-exception clauses can be fetched for the other requirements that the doxastic logic framework delivers but the Lockean–Bayesian framework doesn't. The "[Appendix](#)" contains a list of the requirements that the Lockean–Bayesian and the doxastic logic frameworks have in common, including the modified versions of (R1), (R3), (R4), (R6) and (R8) with their respective exception clauses.

## 9 Concluding remarks

The conclusion from the previous section is that there is a rich kernel of rational requirements of agnosticism that are endorsed by both the Lockean–Bayesian and the doxastic logic frameworks. The large extent to which these alternative frameworks agree is telling. Since these are leading ways of theorizing about the requirements of structural rationality, and since they agree on so much, we can relatively safely make assessments about the coherence of agnostic states by deploying the requirements from that kernel. We get to learn a lot about the requirements for suspended judgment without first resolving the dispute between the two competing accounts.

I submit that the requirements that belong to the aforementioned kernel deliver intuitively correct verdicts about particular cases. For example, suppose I'm a Pyrrhonian skeptic who suspends judgment about whether *there are material objects*. To the extent that my action of steering away from a moving car manifests my belief that *the car is approaching me*, I count as incoherent according to those requirements (seeing as I'm aware that if a car is approaching me then there are material objects). In particular, I am failing to abide by (R7).<sup>33</sup> And, pre-theoretically, there does seem to be a conflict in my attitudes.<sup>34</sup> Or suppose that I believe that *physicalism is true if and only if zombies are not possible*. I am also a physicalist. But when I try to conceive of a scenario featuring zombies in it, I just can't decide whether this is possible or not, and so I suspend judgment about whether *zombies are possible* (I took this example from Rosenkranz 2007). Surely I'm incoherent here as well. Accordingly, (R0) will tell us that I am rationally required not to be in that state.<sup>35</sup>

To get a broader picture of what that kernel of requirements amounts to, one could assemble all of them together in a single system of requirements and work out their joint implications. Furthermore, one could expand that kernel by adding other purported requirements of agnosticism to it—e.g., the 'metacoherence' requirement that one does not believe that  $p$  while suspending judgment about whether one knows that  $p$  (Huemer 2011). Acceptance or rejection of such principles might then depend on (among other things) how much agreement with our pre-theoretic judgments is preserved after we add them to our kernel of rational requirements.

Naturally, there are other formal systems from which one might attempt to derive coherence requirements for belief and suspended judgment (other than the probability calculus and standard doxastic logic). One might try using some non-

<sup>33</sup> (R7) One is rationally required to be such that: if one believes that  $(p \supset q)$ , suspends judgment about  $q$  and does not believe that  $\neg p$ , then one suspends judgment about  $p$ .

<sup>34</sup> Of course, the contextualist might try to 'save my rationality' by pointing out that my attitudes are sensitive to context. But we can add more details to the example in order to fix a single context for my attitudes, thus making it clearer that I am in a conflicting state.

<sup>35</sup> (R0) One is rationally required to be such that: if one believes that  $(p \equiv q)$  and one suspends judgment about whether  $p$ , then one suspends judgment about whether  $q$ .

standard system of doxastic/epistemic logic,<sup>36</sup> or maybe the framework of *imprecise probabilities*.<sup>37</sup> In this first investigation into the matter, however, I have opted for the more canonical systems which have been used to theorize about structural rationality. What purported requirements of rationality will be derived from these alternative systems—and how they fare in comparison to the ones that belong to this paper’s kernel of requirements—is something that is left for future investigation.

**Acknowledgements** Open Access funding provided by Projekt DEAL. I’m indebted to Eyal Tal, Sofia Bokros and Jane Friedman for their comments on a previous version of this paper. I would also like to thank the Alexander von Humboldt Foundation and the German Federal Ministry of Education and Research for funding this research.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## Appendix

### Fundamental requirements of agnosticism, as vindicated by the Lockean–Bayesian (from Sect. 2)

Here are the informal proofs of the basic requirements (SS), (SB) and (SE) using the Lockean–Bayesian framework.

<sup>36</sup> See Rosenkranz (2018) and Berto and Hawke (*forthcoming*) for two recent uses of non-standard systems of epistemic logic.

<sup>37</sup> In this framework, a credal state is represented by a *set* of probability functions. See Sturgeon (2010) for the idea of suspension as ‘thick confidence’, which is formally represented as imprecise probability. The more agnostic or undecided I am about whether  $p$  is the case, the thicker is the interval of probabilities that represents my attitude toward  $p$ . It might be thought that, where  $p, q$  are contingent propositions such that I have no evidence whatsoever pro/con either of them, my agnosticism about these propositions should be represented by a unit interval of probabilities, i.e.  $[Pr_0(p), Pr_1(p)] = [Pr_0(q), Pr_1(q)] = [0, 1]$ . But suppose that  $p$  entails  $q$  and not vice-versa. Intuitively, I should be *less* agnostic about  $q$  than I am about  $p$ , if only a tiny bit—after all,  $q$  has a ‘higher chance’ of being true than  $p$  does. So all functions from  $Pr_0$  to  $Pr_1$  should be such that  $Pr_i(q) > Pr_i(p)$  and, therefore, it is impossible that  $[Pr_0(p), Pr_1(p)] = [Pr_0(q), Pr_1(q)] = [0, 1]$ . See Schoenfield (2012) and also Rinard (2013) for this important point. But then how else is my agnosticism about  $p$  and  $q$  supposed to be represented here? After all, I am totally in the dark as to whether  $p$  and  $q$  are true, and so any subinterval of  $[0, 1]$  would appear to be arbitrary.

- (SS) One is rationally required to be such that: one suspends judgment about  $p$  if and only if one suspends judgment about  $\neg p$ .

*Informal proof:* By the probability calculus,  $Pr(\neg p) = 1 - Pr(p)$ . Now assume that  $Pr(p) = m \in [1 - t, t]$ . So  $Pr(\neg p) = 1 - m$ . But  $m$  can only be as high as  $t$ ; therefore,  $1 - m$  can only be as low as  $1 - t$ . And  $m$  can only be as low as  $1 - t$ ; therefore,  $1 - m$  can only be as high as  $t$ . And, given that  $m \in [1 - t, t]$ , it follows that  $1 - m \in [1 - t, t]$ . Similarly, given that  $1 - m \in [1 - t, t]$ , it follows that  $m \in [1 - t, t]$ . Therefore,  $m \in [1 - t, t]$  iff  $1 - m \in [1 - t, t]$ . So  $Pr(p)$  is middling iff  $Pr(\neg p)$  is middling, for any probability function  $Pr$ . According to the Lockean–Bayesian framework, our credal states are required to be probabilistically coherent—so (SS) follows.

- (SB) One is rationally required to be such that: if one suspends judgment about  $p$  then one does not believe that  $p$ .

*Informal proof:*  $Pr$  is a function. So it is never the case that  $Pr(p) \in [1 - t, t]$  and  $Pr(p) > t$ . Given (L<sub>S</sub>) and (L<sub>B</sub>) and probabilism, it follows that one is never permitted to suspend judgment about whether  $p$  and believe that  $p$  at the same time. Therefore, (SB).

- (SE) One is rationally required to be such that: if one suspends judgment about  $p$  and one has maximal credence that  $(p \equiv q)$ , then one suspends judgment about  $q$ .

*Informal proof:* Let  $Pr(p \equiv q) = 1$ . By the probability calculus,  $Pr(p \supset q) = Pr(q \supset p) = 1$ . So it cannot be that  $Pr(q) < Pr(p)$ , and it cannot be that  $Pr(p) < Pr(q)$ . Therefore,  $Pr(p) \in [1 - t, t]$  iff  $Pr(q) \in [1 - t, t]$ . Given (L<sub>S</sub>), (L<sub>B</sub>) and probabilism, then, (SE) follows.

**Entailment relations that hold in any normal system of doxastic logic at least as strong as  $D$  (from Sect. 6)**

It is easy to prove (D1)–(D0) in any modal logic system at least as strong as  $D$ . I prove (D1) and (D6) as an illustration.

- (D1)  $Bp, Sq \models S(p \wedge q)$

*Informal proof:* Our definition of the  $S$  operator is (Def)  $Sp =_{def}. (\neg Bp \wedge \neg B\neg p)$ . So (D1) boils down to  $Bp, \neg Bq \wedge \neg B\neg q \models \neg B(p \wedge q) \wedge \neg B\neg(p \wedge q)$ . Now assume that  $M, w \models Bp$  and also that  $M, w \models \neg Bq \wedge \neg B\neg q$ , for some random model  $M = \langle W, R, V \rangle$  where  $R$  is serial and  $w \in W$ . Suppose that (a)  $M, w \models B(p \wedge q)$ . But this would require that  $M, w \models Bq$ , and it therefore contradicts our initial assumption that  $M, w \models \neg Bq \wedge \neg B\neg q$ . So  $M, w \models \neg B(p \wedge q)$ , contrary to supposition (a). Suppose now that (b)  $M, w \models B\neg(p \wedge q)$ . Our initial supposition that  $M, w \models Bp$  entails, together with (b), that  $M, w \models B\neg q$ . But this also contradicts our initial supposition that  $M, w \models \neg Bq \wedge \neg B\neg q$ . So  $M, w \models \neg B\neg(p \wedge q)$ , contrary to (b). So  $M, w \models \neg B(p \wedge q) \wedge \neg B\neg(p \wedge q)$ , for any such  $M$  and  $w$  with  $M, w \models Bp$  and  $M, w \models \neg Bq \wedge \neg B\neg q$ —that is, (D1).

$$(D6) \quad S(p \vee q) \models Sp \vee Sq$$

*Informal proof:* (D6) boils down to  $\neg B(p \vee q) \wedge \neg B\neg(p \vee q) \models (\neg Bp \wedge \neg B\neg p) \vee (\neg Bq \wedge \neg B\neg q)$ . Assume that  $M, w \models \neg B(p \vee q) \wedge \neg B\neg(p \vee q)$ , for some random model  $M$ . Now suppose that (a)  $M, w \models \neg((\neg Bp \wedge \neg B\neg p) \vee (\neg Bq \wedge \neg B\neg q))$ . It follows that  $M, w \models \neg(\neg Bp \wedge \neg B\neg p) \wedge \neg(\neg Bq \wedge \neg B\neg q)$  (DeMorgan's law). So  $M, w \models \neg(\neg Bp \wedge \neg B\neg p)$ , ergo  $M, w \models Bp \vee B\neg p$ . From the initial supposition that  $M, w \models \neg B(p \vee q) \wedge \neg B\neg(p \vee q)$ , however, it follows that  $M, w \models \neg Bp$ . But from  $M, w \models Bp \vee B\neg p$  and  $M, w \models \neg Bp$  it follows that  $M, w \models B\neg p$ . A similar line of reasoning leads us to conclude that  $M, w \models B\neg q$  on the basis of (a). These two results together, however, give us  $M, w \models B\neg p \wedge B\neg q$ , ergo  $M, w \models B\neg(p \vee q)$ , which contradicts our initial assumption. So (a) must be false given our initial assumption, that is, (D6).

### The kernel of requirements that are vindicated by both the Lockean–Bayesian and the doxastic logic frameworks (from Sect. 8)

Here are some requirements that belong to the shared kernel mentioned in the paper. I use  $R$  as the requirement operator to express the relevant requirements. ' $R\Phi$ ' reads as follows: one is required to be such that  $\Phi$ .  $B$  is the belief-operator, and  $S$  the suspended judgment-operator. Where  $\Phi_1, \dots, \Phi_n, \Psi$  are doxastic formulas, the material conditional is always to be interpreted as the main connective of the formula that is under the scope of  $R$  in formulas of type  $R(\Phi_1 \wedge \dots \wedge \Phi_n \supset \Psi)$ .

- (SS)  $R(Sp \equiv S\neg p)$
- (SB)  $R(Sp \supset \neg Bp)$
- (R1\*)  $R(Bp \wedge Sq \wedge \neg B\neg(p \wedge q) \supset S(p \wedge q))$
- (R2)  $R(Sp \wedge Sq \wedge \neg B\neg(p \wedge q) \supset S(p \wedge q))$
- (R3\*)  $R(S(p \wedge q) \wedge Bp \wedge \neg Bq \supset Sq)$
- (R4\*)  $R(B\neg p \wedge Sq \wedge \neg B(p \vee q) \supset S(p \vee q))$
- (R5)  $R(Sp \wedge Sq \wedge \neg B(p \vee q) \supset S(p \vee q))$
- (R6\*)  $R(S(p \vee q) \wedge \neg B\neg p \wedge \neg B\neg q \supset Sp \vee Sq)$
- (R7)  $R(B(p \supset q) \wedge Sq \wedge \neg B\neg p \supset Sp)$
- (R8\*)  $R(Bp \wedge Sq \wedge \neg B(p \supset q) \supset S(p \supset q))$
- (R9)  $R(S(p \supset q) \wedge S(p \supset \neg q) \supset Sq)$
- (R0)  $R(B(p \equiv q) \wedge Sp \supset Sq)$ .

### References

- Berto, F., & Hawke, P. (forthcoming). Knowability relative to information. *Mind*.  
 Broome, J. (1999). Normative requirements. *Ratio*, 12(4), 398–419.  
 Broome, J. (2007). Wide or narrow scope. *Mind*, 116(462), 359–370.  
 Brunero, J. (2010). The scope of rational requirements. *The Philosophical Quarterly*, 60(238), 28–49.  
 Buchak, L. (2014). Belief, credence and norms. *Philosophical Studies*, 169(2), 285–311.  
 Christensen, D. (2004). *Putting logic in its place*. Oxford: Oxford University Press.  
 Clarke, R. (2013). Belief is credence one (in context). *Philosopher's Imprint*, 13, 1–18.

- Dodd, D. (2016). Belief and certainty. *Synthese*, 194(11), 4597–4621.
- Douven, I., & Williamson, T. (2006). Generalizing the lottery paradox. *The British Journal for the Philosophy of Science*, 57(4), 755–779.
- Easwaran, K., & Fitelson, B. (2015). Accuracy, coherence, and evidence. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 5, pp. 61–96). Oxford: Oxford University Press.
- Foley, R. (1993). *Working without a net*. Oxford: Oxford University Press.
- Friedman, J. (2013a). Suspended judgment. *Philosophical Studies*, 162(2), 165–181.
- Friedman, J. (2013b). Rational agnosticism and degrees of belief. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 4, pp. 57–81). Oxford: OUP.
- Friedman, J. (2017). Why suspend judging? *Noûs*, 51(2), 302–326.
- Greco, D. (2015). How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29(1), 179–201.
- Hawthorne, J., & Bovens, L. (1999). The preface, the lottery, and the logic of belief. *Mind*, 108, 241–264.
- Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Cornell: Cornell University Press.
- Huemer, M. (2011). The puzzle of metacoherence. *Philosophy and Phenomenological Research*, 82(1), 1–21.
- Hughes, G. E., & Cresswell, M. (1996). *A new introduction to modal logic*. London: Routledge.
- Jackson, E. (2019). Belief and credence: Why the attitude-type matters. *Philosophical Studies*, 176, 2477–2496.
- Joyce, J. M. (1998). A non-pragmatic vindication of probabilism. *Philosophy of Science*, 65(4), 575–603.
- Kiesewetter, B. (2017). *The normativity of rationality*. Oxford: Oxford University Press.
- Kolodny, N. (2007). How does coherence matter? *Proceedings of the Aristotelian Society*, 107(1), 229–263.
- Kyburg, H. (1970). *Probability and inductive logic*. Toronto: MacMillan.
- Leitgeb, H. (2014). The stability theory of belief. *Philosophical Review*, 123(2), 131–171.
- Levi, I. (1991). *The fixation of belief and its undoing: Changing beliefs through inquiry*. Cambridge: Cambridge University Press.
- Littlejohn, C. (2015). Who cares what you accurately believes? *Philosophical Perspectives*, 29(1), 217–248.
- MacFarlane, J. (2004). In what sense (if any) is logic normative for thought.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25(6), 205–207.
- Nelkin, D. (2000). The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109, 373–409.
- Pettigrew, R. (2016). *Accuracy and the laws of credence*. Oxford: Oxford University Press.
- Pollock, J. (1986). The paradox of the preface. *Philosophy of Science*, 53, 246–258.
- Pollock, J. (1990). *Nomic probability and the foundations of induction*. Oxford: Oxford University Press.
- Rinard, S. (2013). Against radical credal imprecision. *Thought*, 2(1), 157–165.
- Rosenkranz, S. (2007). Agnosticism as a third stance. *Mind*, 116(461), 55–104.
- Rosenkranz, S. (2018). The structure of justification. *Mind*, 127(506), 309–338.
- Ross, J., & Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2), 259–288.
- Ryan, S. (1991). The preface paradox. *Philosophical Studies*, 64, 293–307.
- Schoenfield, M. (2012). Chilling out on epistemic rationality. *Philosophical Studies*, 158(2), 197–219.
- Smithies, D. (2015). Ideal rationality and logical omniscience. *Synthese*, 192(9), 2769–2793.
- Staffel, J. (2016). Beliefs, buses and lotteries. *Philosophical Studies*, 173(7), 1721–1734.
- Stalnaker, R. (2006). On logics of knowledge and belief. *Philosophical Studies*, 128(1), 169–199.
- Steinberger, F. (2016). The normative status of logic. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2017 edition). Stanford: Stanford University.
- Sturgeon, S. (2010). Confidence and coarse-grained attitudes. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3). Oxford: Oxford University Press.
- Titelbaum, M. (2015). Rationality's fixed point (or: In defense of right reason). In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 5). Oxford: Oxford University Press.
- Way, J. (2011). The symmetry of rational requirements. *Philosophical Studies*, 155(2), 227–239.
- Way, J. (2018). Reasons and rationality. In D. Star (Ed.), *The Oxford handbook of reasons and normativity*. Oxford: Oxford University Press.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 36(16), 267–297.

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

Yap, A. (2014). Idealization, epistemic logic and epistemology. *Synthese*, 191(14), 3351–3366.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.