



# Ought-contextualism and reasoning

Darren Bradley<sup>1</sup> 

Received: 29 January 2020 / Accepted: 14 October 2020 / Published online: 19 January 2021  
© The Author(s) 2021

## Abstract

What does logic tell us about how we ought to reason? If  $P$  entails  $Q$ , and I believe  $P$ , should I believe  $Q$ ? I will argue that we should embed the issue in an independently motivated contextualist semantics for ‘ought’, with parameters for a standard and set of propositions. With the contextualist machinery in hand, we can defend a strong principle expressing how agents ought to reason while accommodating conflicting intuitions. I then show how our judgments about blame and guidance can be handled by this machinery.

**Keywords** Ought-contextualism · Reasoning · Inferences

## 1 Introduction

What does logic tell us about how we ought to reason? If  $P$  entails  $Q$ , and you believe  $P$ , should you believe  $Q$ ? There seem to be cases where you should not, for example, if you have evidence against  $Q$ , or the inference is not worth making. So we need a theory telling us when an inference ought to be made, and when not. I will argue that we should embed the issue in an independently motivated contextualist semantics for ‘ought’. With the contextualist machinery in hand we can give a theory of when inferences should be made and when not.

Section 2 explains the background and the main problems connecting logic with norms of reasoning. Section 3 explains the two parameters we need for contextualism about ‘ought’—a set of live possibilities and a standard. Section 4 discusses the objection from belief revision (this and the other problems will be explained in Sect. 2) and argues that it can be solved by using the set of live possibilities, as can the preface paradox (Sect. 5) and the problem of excessive demands (Sect. 6). Section 7 discusses the problem of clutter avoidance and argues that it can be solved by using the rele-

---

✉ Darren Bradley  
bradleydarren@gmail.com

<sup>1</sup> Philosophy Department, Leeds University, Woodhouse Lane, Leeds LS2 9JT, UK

vant standard. Section 8 discusses the implications for blame and guidance. Section 9 concludes.

## 2 Background

What is the relation between logic and reasoning? For example, suppose an agent believes that P. Suppose also that Q is a logical consequence of P, but leave open whether the agent believes that Q is a logical consequence of P.<sup>1</sup> Should the agent infer that Q? (To fill out the example, P might be ‘it’s raining and if it’s raining then it’s wet’ and Q might be ‘it’s wet’.)

A useful starting point is that logic ‘prescribe[s] universally how one ought to think’ (Frege 1893/1903/2009, p. 15). This suggests that the agent ought to believe Q. We might try to capture the idea with the following norm:<sup>2</sup>

### *Strong Normativity Thesis*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought to believe Q

(The reason for calling it ‘strong’ will emerge just below).<sup>3</sup>

Various purported counter-examples have been given (Harman 1986, Field 2009, 2015, MacFarlane ms) based around four main problems:<sup>4</sup>

### *Belief Revision*

The agent might have strong evidence against Q. If so, they should surely revise their belief that P, rather than believe Q.<sup>5</sup>

### *The Preface Paradox*

Suppose S rationally believes each of the assertions in his book, P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub>. Let Q stand for the conjunction, P<sub>1</sub> & P<sub>2</sub>...&P<sub>n</sub>. Q is entailed by the author’s beliefs. But surely, since the author regards himself as fallible, he should not believe the conjunction of all his assertions.<sup>6</sup>

<sup>1</sup> I follow the literature in assuming that logical consequence is not dependent on thought, reasoning, or minds (see Prawitz 2005). Other than that, any variety of logical consequence can be plugged in. See Steinberger (2019b) for discussion. Modus ponens is the standard example, but what I say is intended to generalize to other deductive principles, and, *mutatis mutandis*, to inductive principles.

<sup>2</sup> I take conditionals like this to be material conditionals. It is a narrow-scope requirement.

<sup>3</sup> Part of the reason is to distinguish it from Steinberger’s (2019a) similar but not obviously identical ‘Normativity Thesis’.

<sup>4</sup> One problem I won’t discuss is the explosion caused by inconsistent beliefs (Allo 2016, Steinberger 2016). I don’t think contextualism can solve this problem. I suspect the solution is to limit the antecedent of the norm to mental states more basic than beliefs, such as experiences, but that is a topic for another paper. Also, the problems above are based on cases where the agent fails to make a valid inference. There are mirror-image cases where agents *do* make inferences which are *invalid* e.g. Pascal’s wager. I leave the extension of this framework to such cases for future work.

<sup>5</sup> I take this to be the same as the ‘bootstrapping problem’ (see Broome 2013, Sect. 5.3; Gibbons 2013, p. 32), or at least has the same solution. .

<sup>6</sup> See Makinson (1965). This is similar to the lottery paradox (Kyburg 1961) but I leave a discussion of the lottery paradox for another occasion.

### *Excessive Demands*

Some consequences of an agent's beliefs are too complicated for them to work out. For example, Fermat's Last Theorem follows from the rules of arithmetic. But surely most humans who know the rules of arithmetic have no obligation to believe Fermat's Last Theorem.

### *Clutter Avoidance*

Some consequences of an agent's belief are too uninteresting to be worth working out. For example, an agent might be able to infer that either grass is green or Elvis lives on the moon, using disjunction introduction. But surely they have no obligation to make such an inference, and might be irrational for doing so.

One reaction to these problems is to weaken the *Strong Normativity Thesis*. To see how this might be done, it is helpful to review MacFarlane's taxonomy of three choice-points:

1. Type of deontic operator. Do facts about logical validity give rise to strict obligations, permissions, or reasons for belief? Does the agent have
  - i. a requirement/obligation to believe Q
  - ii. permission to believe Q
  - iii. a reason to believe Q?
2. Polarity. Are these obligations/permissions/reasons to
  - i. believe
  - ii. not to disbelieve?
3. Scope of deontic operator. Does the deontic operator govern:
  - i. the consequent of the conditional (if S believes P then S ought to believe Q)
  - ii. both the antecedent and the consequent (if S ought to believe P then S ought to believe Q)
  - iii. the whole conditional (S ought to believe that if P then Q)?

For each choice-point, the options described move roughly<sup>7</sup> from more demanding to less demanding. For example, a norm that says agents are required to believe Q is more demanding than a norm that says agents merely have permission to believe Q. The *Strong Normativity Thesis* takes the first, and most demanding, on all three choice-points. It says that agents have a requirement rather than a permission or reason, that the requirement is to believe rather than merely not disbelieve (which includes suspension of belief), and that the requirement attaches to only the consequent.<sup>8</sup> The existing literature largely considers weakening the *Strong Normativity Thesis* by moving to these other choice-points.<sup>9</sup> But I don't think these choice points are the right places to weaken the link between logic and reasoning.

<sup>7</sup> I add 'roughly' because it's not obvious that having permission is stronger than having a reason.

<sup>8</sup> See Schroeder (2004), Kolodny (2005) and Finlay (2014, pp. 52–53) for arguments against the wide-scope solution. See Titelbaum (2015) for references and discussion of the logical relations between wide and narrow scope norms.

<sup>9</sup> See Steinberger (forthcoming a, c) for discussions based on these choice-points.

I think we need a different way to weaken the *Strong Normativity Thesis*—the key is that ‘ought’ is context-sensitive, and the *Strong Normativity Thesis* is true only with a particular sense of ‘ought’. I will explain this in the next section, then show how the counter-examples are avoided.

Eight quick clarifications (impatient readers can skip to the next section): First, I take reasoning to be a process of transitioning between beliefs. Beyond that I remain neutral on what reasoning is.<sup>10</sup>

Second, I remain neutral on the existence and nature of other epistemic norms in the area e.g. norms of belief (Fassio 2019), norms that agents should collect more evidence (Friedman forthcoming) etc.

Third, I remain neutral on whether norms of reasoning are fundamental or derived from more fundamental synchronic norms (Hedden 2015).

Fourth, I will focus on deductive inferences rather than inductive inferences. I think my account can be extended to inductive inferences, but do not do so here.

Fifth, I take ‘reasoning’ to mean the same as ‘inferring’. The latter is useful for talking about individual inferences, which is a more natural locution than ‘an individual act of reasoning’.

Sixth, I will use ‘correct’, ‘bad’ and ‘good’ only as normative terms, and use ‘valid’ for logical relations.

Seventh, I will assume that there can be epistemic reason to believe, and to make (or not to make) an inference to a belief.<sup>11</sup> I will assume that an inference is a type of action, so there can also be practical reason to make (or not to make) an inference. I will remain neutral on whether there can be practical reasons to believe and on whether epistemic reasons are fundamental or are ultimately grounded in practical reason.<sup>12</sup>

Finally, I take there to be a close connection between ‘ought’, ‘should’, ‘good’ and ‘reasons’. Specifically, I assume that ‘what one *should* do’ is synonymous with ‘what one *ought* to do’, ‘what one has most *reason* to do’ and ‘what is good’ (Shafer-Landau 2005; Broome 2013; Berker 2018). I will focus on contextualism about ‘ought’, but I take this to have straight-forward implications for contextualism about other normative terms (Finlay 2014).<sup>13</sup> I remain neutral on which if any is fundamental.

### 3 Two parameters

Suppose Napoleon, an eighteen century general, and Heimson, a twentieth century schizophrenic, utter the same sentence: ‘I am Napoleon’. There is a sense in which ‘I

<sup>10</sup> Valaris (2019) argues that ‘reasoning’ is ambiguous between cases where the agent believes the propositions and cases where the agent is working out the consequences of propositions under supposition. I follow the literature in talking about cases where the agent believes the propositions, but my arguments apply to cases where the propositions are merely supposed.

<sup>11</sup> See Singer and Aronowitz (forthcoming) for the view that there can be epistemic reason to do pretty much anything.

<sup>12</sup> See Cowie (2014), Woods and Maguire (forthcoming).

<sup>13</sup> My arguments do not need contextualism, they just need some way of putting the parameters into the sentence. I use the familiar contextualist theory, but other semantics theories might work (see Vetter and Viebahn 2016).

means the same thing in both utterances. This type of meaning can be thought of as a rule picking out whoever is speaking; this is the *character* (Perry 1979; Kaplan 1989). And there is a sense in which ‘I’ means different things in each utterance, Napoleon and Heimson respectively; this is the *content*. So the content of any utterance of ‘I’ depends on a parameter—the speaker. We can make the parameter explicit by adding to the text who ‘I’ is relative to e.g. ‘I-Napoleon’ or ‘I-Heimson’.

An analogous view regarding ‘ought’ has become increasingly popular.<sup>14</sup> In fact it is plausible that there are at least two<sup>15</sup> parameters needed to fix the content of a sentence including ‘ought’—a standard and a set of live possibilities. In this section I will explain the view, and also separate the core commitments from stronger positions we need not be committed to.

### 3.1 Propositions/possible worlds

The first parameter is a modal base which determines a proposition or set of live possible worlds.<sup>16</sup> The live worlds are those compatible with the modal base. If the modal base is empty then all worlds are live. As the modal base grows, the set of live worlds is restricted. On the standard theory of modals (Kratzer 1981), ‘it must be that p’ means, roughly, that in all the live worlds, p.

This parameter is often called the ‘information set’, but using *information* here is too restrictive, for two reasons. The first reason is that information is naturally taken to imply truth. However, it will be important that agents can make good inferences from false beliefs.<sup>17</sup> We can allow that the modal base consists of the beliefs of the subject of the sentence, or the speaker, or some third party, or the collective beliefs of a group, or the propositions known by any of the former, or any of these plus a number of fixed propositions, and endless further options.

The second reason is novel. I think information is too restrictive because the parameter needs to vary with the agent’s *abilities*, not just their information. What one ought to do depends on what possibilities you can bring about in the future.<sup>18</sup> To motivate this, suppose you are on the beach and see someone struggling in the water. Whether you ought to dive in depends on whether you can swim. ‘You ought to dive in given that you can swim’ is true while ‘you ought to dive in given that you cannot swim’ is false. In a context in which you can swim, performing the rescue yourself is a live


<sup>14</sup> See Wedgwood (2006, 2007: ch. 5, 2016); Brogaard (2008); Kolodny and MacFarlane (2010), Björnsson and Finlay (2010); Dowell (2012, 2013), Charlow (2013), Cariani et al. (2013), Carr (2015), Chrisman (2015), Silk (2017), Khoo and Knobe (2018), Worsnip (2019a). This work is strongly influenced by Kratzer (1981, 1991, 2012). For earlier forerunners of contextualism in metaethics see Geach (1956), Foot (1972), Harman (1975, 1996) and Dreier (1990).

<sup>15</sup> Finlay (2014, 2016) argues that only a single parameter is needed. For convenience I will help myself to the more profligate and familiar theory.

<sup>16</sup> I will take for granted that propositions can be modelled by possible worlds, and use propositions and possible worlds interchangeably.

<sup>17</sup> The standard semantics assumed factivity. Finlay (2014) uses ‘information’ but without intending factivity (personal communication). See Nebel (2019) for a discussion of the factivity of reasons. .

<sup>18</sup> For an ‘ought’ relativized to the beliefs of the subject, what they ought to do depends on, not just the possible ways they *believe* the world might be in the present and past, but what possibilities they *believe* they can bring about in the future.



**Fig. 1** Live possibilities vary with abilities

possibility; in a context in which you cannot swim, this possibility is ruled out. So the live possibilities can be determined in part by a *set of actions*. With the assumption that making an inference is a type of action,<sup>19</sup> the value of the live possibilities parameter can depend in part on which inferences the agent can make (Fig. 1).

We will also make the standard assumption that when ‘ought’ occurs in the consequent of a conditional, the antecedent of the conditional is added to the modal base. Consider ‘if S believes P then S ought to believe Q’. The ‘ought’ has a modal base which includes ‘S believes P’. We can make this explicit e.g. ‘if S believes P then S ought-given-S-believes-P to believe Q’.

### 3.2 Standard

The second parameter is a standard or goal which determines an ordering of the live possible worlds. Plausibly, ‘S ought to A’ is true iff S A’s in every live world at the top of the ranking.<sup>20</sup>

The standard need not be one that the subject cares about.<sup>21</sup> If I say ‘you ought to start with the cutlery on the outer edge’, the standard might be the rules of etiquette. The more explicit sentence is ‘by standards of etiquette, you ought to start with the cutlery on the outer edge’. This sentence can remain true even if you don’t care about etiquette. This allows us to say to the psychopath ‘you shouldn’t kill people’; the full sentence is ‘by standards of morality, you shouldn’t kill people’, and this is true even if the psychopath doesn’t care about morality.<sup>22</sup>

For our purposes we only need to distinguish two standards: those corresponding to the *epistemic ought* and the *practical ought*.<sup>23</sup> We can get a grip on the *epistemic ought*

<sup>19</sup> Some disagree. For example, Ram Neta holds that ‘every inference is simply a judgment with a certain kind of content’ (2013: 404). If so, the second reason for including actions among the possibilities does not apply.

<sup>20</sup> Kratzer’s (1981) semantics provides an ordering of worlds. However, the ordering of worlds can be changed by a change in information (Kolodny and MacFarlane 2010). So I prefer to use a standard, which is plausibly more fundamental than the ordering.

<sup>21</sup> This corresponds to a normative rather than a teleological interpretation of the associated norm (Shah 2003; Fassio 2019).

<sup>22</sup> See Sinnott-Armstrong (2014).

<sup>23</sup> I will assume for simplicity that each is associated with a unique standard.

by thinking about contexts where the conversation concerns some epistemic standard such as having true beliefs. Typical sentences might be ‘you ought to be uncertain’ or ‘we ought to expect defeat’.

Again, the standard need not be one that the subject cares about, so we need not assume that agents care about any epistemic goals. For example, someone who is told how a film ends ought (in the epistemic sense) to believe what they are told, even if they do not care how it ends, and even if they don’t want to know how it ends.<sup>24</sup> The full sentence might be ‘by epistemic standards, you should believe that this is how the film ends’.

There is disagreement about what the epistemic standard is. Leading contenders for epistemic standards include having beliefs that are (a) true (b) justified (c) knowledge.<sup>25,26</sup> The differences between these positions won’t matter here, so I will remain neutral. And we can remain neutral on whether the standard (e.g. truth) is constitutive of belief or whether something can be a belief without having such a standard.<sup>27</sup>

This brings us to the *practical* ought.<sup>28</sup> We can get a grip on the *practical* ought by thinking about normal contexts where the conversation concerns what is best to do. Typical sentences might be ‘you ought to stay in school’ or ‘should I boil or steam the vegetables?’. Call the standard associated with the practical ought the *practical standard*.

There is disagreement about what the practical standard is. Humeans hold that the practical standard is a function of one’s desires e.g. the standard might be to maximize a weighted set of desires. Non-Humeans might hold that the practical standard is to maximize value. There are further debates about whether the practical standard is to maximize actual value or expected value, and whether expected value is determined by beliefs or evidence.<sup>29</sup> We can remain neutral on these issues.<sup>30</sup>

We can also remain neutral on whether there are further parameters which determine the content beyond standards and propositions. For example, Carr (2015) argues that ‘ought’ must be relativized to a decision rule. This may be so, but it will not play a role below.

<sup>24</sup> See Kelly (2003).

<sup>25</sup> See Chignell (2018).

<sup>26</sup> I set aside epistemic norm pluralism (Hughes 2017; Kopec 2018).

<sup>27</sup> See Wedgwood (2002), Boghossian (2003)

<sup>28</sup> See Williams (1965), Harman (1973) and Geach (1982), Broome (2013, pp. 12–24), and Schroeder (2011), Kiesewetter (2017, Sect. 1.2.3).

<sup>29</sup> See Kiesewetter (2017, ch. 8) for a helpful overview and further references.

<sup>30</sup> We can also remain neutral on whether there is a metaphysically privileged normative ‘ought’ which expresses genuine normative authority. Worsnip writes:

“we should be careful to separate the question of whether (e.g.) the law ...has genuine normative authority from whether there is a robustly normative usage of the legal ‘ought’. The former requires the law to actually possess normative authority, whereas the latter only requires there to be speakers who take the law to possess normative authority. So even if only a handful of the above ‘oughts’ reflect a genuine source of normativity, many more of them might nevertheless be robustly normative usages of ‘ought’.” (Worsnip 2019a, page numbers not yet available; see also Worsnip 2020)

The former question is metaphysical, the latter is semantic. We need only make the latter assumption that there are robustly normative usages of ‘ought’.

Now we have this machinery on the table, I will argue that the problems regarding the norms of reasoning can be resolved. There are numerous precisifications of the Strong Normativity Thesis, some of which are true and some of which are false.

#### 4 Objection from belief revision

Suppose S believes P, believes P entails Q, but S has strong evidence against Q. It seems that S should not come to believe Q. But this is difficult to accommodate if our principle has a claim in the consequent about what the agent should believe e.g.

##### *Strong Normativity Thesis*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought to believe Q

This is the problem of belief revision.<sup>31</sup>

To solve this problem (and the next) we need to distinguish the normative status of a *belief* from the normative status of an *inference*. Crucially, an agent can make a good inference from a bad (e.g. unjustified) belief.<sup>32</sup> Our question concerns reasoning, so we want to bracket the question of whether the initial belief was justified and focus on the question of whether the inference was good. So the solution to the problem of belief revision is to say that the inference to Q is good but the belief that Q is not.

What role does contextualism play here? It helps specify a modal base relative to which the inference is good. So we modify the *Strong Normativity Thesis* in two ways—we replace ‘believe’ with ‘infer’ and we make explicit the modal base:

##### *Modified Strong Normativity Thesis*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought-given-P-entails-Q-and-S-believes-P to infer Q

This allows us to judge that the inference to Q is good *qua* inference, while remaining neutral on the epistemic status of the initial belief that P, and consequently remaining neutral on the epistemic status of a belief that Q.

Someone might object that it is overall justification for the *belief* that Q that we are really interested in. If so, an account that brackets the rest of an agent’s epistemic states is unhelpful.

The first response is to flat-footedly reply that our question is about reasoning, not belief. But even if our concern were about overall justification of belief, our framework would be helpful—we would just have to identify some initial epistemic state (e.g. basic beliefs, or evidence<sup>33</sup>), then plug in the agent’s total initial epistemic state for P:

<sup>31</sup> One response is to weaken the consequent from the obligation operator to the permissible operator (Broome 2013, p. 219). Another is to endorse wide scope norms (Broome 1999).

<sup>32</sup> See Broome (1999, pp. 418–419) ‘In your reasoning, you can take as premises beliefs and intentions you have no reason to have, and even beliefs and intentions you ought not to have. The nature of your reasoning is unaffected by whether or not you ought to have the beliefs and intentions it is premised on.’

<sup>33</sup> For basic beliefs see BonJour (1985), for evidence see Feldman and Conee (1985).



*Modified Strong Total Normativity Thesis*

For all agents S, and propositions P and Q:

If S's initial-epistemic-state entails Q, then S ought-given-S's-initial-epistemic-state to believe Q

So this reasoning framework can be placed into a bigger story about rational belief. But rational belief raises numerous tricky issues such as internalism, defeaters and inductive reasoning which go beyond the scope of this paper.

Terminology: In future, rather than repeating the whole antecedent appended to 'ought', I will just write 'ought<sub>A</sub>'.

## 5 The preface paradox

In this section I will argue that the same response, that of relativizing 'ought' to a possibilities parameter, solves the preface paradox:

*Preface Paradox*

Suppose S rationally believes each of the assertions in his book, P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub>. Let Q stand for the conjunction, P<sub>1</sub> & P<sub>2</sub>...&P<sub>n</sub>. Since the author regards himself as fallible, he should not believe the conjunction of all his assertions (Q).

Thus S believes P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub> and that they entail Q, but S should not believe Q. The problem is usually taken to be that of explaining why the author should not make the inference to Q.

But there is a sense in which the author *should* make the inference. If we move from talk of belief to talk of inferences, and set the live possibilities parameter to the proposition that S rationally believes P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub>, then we can hold that the inference to Q is correct after all:

*Modified Strong Normativity Thesis*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought<sub>A</sub> to infer Q

So there is a sense in which S ought to infer Q.

Someone might object that this misses the point, arguing that the Preface Paradox shows that the agent should not make the inference. But why not? The standard answer begins with the observation that each of P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub> has partial justification, and then concludes that justification for Q might fall below some threshold.<sup>34</sup>

But the level of justification of each of P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub> is not at issue. We are assessing the *deductive inference* from P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub> to Q (not the belief that Q) and thereby setting aside the justificatory status of P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub>. It is irrelevant whether each of P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub> has only partial justification, or is even completely unjustified. We are asking whether it is correct to infer Q from the set of beliefs that P<sub>1</sub>, P<sub>2</sub>...P<sub>n</sub>, and indeed it is.

It might be useful to draw an analogy with Lewis's (1980) Principal Principle. Roughly, it says that an agent should have credence of x in P *given* that they are

<sup>34</sup> See Worsnip (2019b) for discussion and references.

rationally certain that the chance of P is x. If the agent is not rationally certain what the chance is, then the Principal Principle says nothing directly about what the agent should believe. Similarly, if the agent does not believe each of P1, P2...Pn then the *Modified Strong Normativity Thesis* says nothing directly about what the agent should believe. In fact the *Modified Strong Normativity Thesis* never says anything directly about what an agent should believe. It just says that given the set of beliefs P1, P2...Pn they should *infer* Q.

## 6 Excessive demands

The problem of excessive demands is that we sometimes cannot work out the consequences of our beliefs. All the theorems of mathematics follow from our beliefs about arithmetic, but surely we are not required to infer them.

This conflict between logic and the norms of reasoning can be resolved by again invoking the live possibilities parameter. Above we used the *belief* part of the live possibilities; here we use the *actions* part of the live possibilities, invoking the assumption that an inference is a type of action.

From any belief there are an infinite number of valid inferences that could be made, of which some are simple and some are complicated. Let's first focus on the infinite set of valid inferences. There is a sense of 'ought' which includes all valid inferences in the possibilities parameter. (Bayesians will be familiar with this, as it is what 'rational' usually means in the Bayesian literature.<sup>35</sup>) We can make this ideally rational 'ought' explicit by using 'ought-rationally'. And we can make explicit the sense of 'ought' which is limited to inferences some particular agent is able to make with 'ought-actually'. We get a false principle if we combine ought-actually with all the valid inferences:

### *False Requirement (FR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought<sub>A</sub>-*actually* to infer Q

This implies that S ought-actually to believe all theorems of mathematics. This is the root of the problem of excessive demands.

But the objection is side-stepped if we use ought-rationally:

### *Rational requirement (RR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S ought<sub>A</sub>-*rationally* to infer Q<sup>36</sup>

<sup>35</sup> For a recent paper on this topic, see Jackson (2018). Do not confuse this idealized ought with that used in defences of the truth norm of belief (see McHugh 2012), which I suggest is the sense of ought relativized to full information in the modal base.

<sup>36</sup> MacFarlane (ms) seems to have this in mind when he writes: 'We seek logical knowledge so that we will know how we ought to revise our beliefs: not just how we will be obligated to revise them when we acquire this logical knowledge, but how we are obligated to revise them even now, in our state of ignorance.' p. 12.

The objection is also side-stepped if we use ought-actually and add to the antecedent that the inferences are those S is able to make, which we can call ‘S-available inferences’.

*Non-rational requirement (NR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-available-inferences, then S ought<sub>A</sub>-actually to infer Q

Thus the live possibilities parameter solves the problem of excessive demands by providing a reading of the Strong Normativity Thesis on which one is not required to infer all the theorems of mathematics (NR), and it also explains the intuition that there is a sense in which you should infer all the theorems of mathematics (RR).

To fill this out, we can imagine three different sentences which can be inferred from P.

Q1: An obvious inference that any reasoner can make

Q2: A difficult inference that a logic student can make

Q3: A superhuman inference that no human can make

Let w3 be the world where S infers Q3, Q2 and Q1; let w2 be the world where S infers Q2 and Q1; and let w1 be the world where S infers only Q1. Worlds are ordered (vertically) by how well they achieve the epistemic goal (Fig. 2).




Fig. 2 Three worlds in which increasingly complex inferences are made

Agents ought to make true the best live world. Relative to ought-rationally, S ought to make w3 true. But if w3 is not live then S ought-actually make only w2 true. And if S is unable to do difficult logical reasoning then only w1 is live.

One complication concerns which inferences the agent is able to make, as there is some flexibility about what is held constant. Suppose the agent is tired, and this restricts the inferences they do make. What should we hold fixed when assessing what inferences they are able to make? If we hold fixed that the agent is tired, we will get one set of available inferences. If we allow them a nap we will get a bigger set of inferences. If we allow them to take a mathematics course, we will get an even bigger set of inferences.

I don't want to take a stand on this, a topic which has been discussed in the literature on ought-implies-can.<sup>37</sup> I suspect that 'able' is also context-sensitive, which would make 'available inferences' context-sensitive. And any vagueness in 'available' will be matched by vagueness in 'ought'. The truth of NR just requires that the demands of ought-actually do not extend beyond the available inferences.<sup>38</sup>

We have explained how excessive demandingness can be avoided by positing the relatively undemanding norm of NR. But we'll see that NR might still be too demanding.

## 7 Clutter avoidance

NR (and RR) still face the problem of clutter avoidance. They seem to imply that I am obligated to believe all of the infinitely many trivial logical consequences of my beliefs. This looks implausible. Steinberger (2019, p. 11) writes:

Not only do I not care about, say, the disjunction 'I am wearing blue socks or Elvis Presley was an alien' entailed by my true belief that I am wearing blue socks, it would be positively irrational for me to squander my meagre cognitive resources on inferring trivial implications of my beliefs that are of no value to my goals.

Many philosophers have concluded that there must be a no-clutter norm,<sup>39</sup> but these cause serious problems.<sup>40</sup>

I think the problem of clutter avoidance can be solved by invoking the parameter of the standard. I will argue that in normal contexts you ought not to make trivial inferences, yet we can identify contexts in which you should make trivial inferences.

The sense of 'ought' in which you ought to infer all the trivial logical consequences of your beliefs is the epistemic sense (e.g. the standard of having all and only true beliefs).<sup>41</sup> We can make this parameter value explicit with 'epistemically-ought':

### *Non-rational Epistemic-Requirement (NER)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-available-inferences, then S *epistemically-ought*<sub>A</sub>-actually to infer Q

<sup>37</sup> See Frankena (1958) and Schwarz (2020) for a contextualist account of 'can'.

<sup>38</sup> Someone might object that this is trivial or analytically true. Yes. We are trying to find true principles which vindicate the intuitions behind plausible but overly strong principles. We should expect some of these true principles to be trivial or analytic.

<sup>39</sup> See Harman (1986) Goldman (1986), Christensen (1994), Williamson (1998), Ryan (1999), Feldman (2000), Wallace (2001), Sainsbury (2002), DePaul (2004), MacFarlane (2004), White (2005), Field (2009), Douven (2010) and Wedgwood (2012).

<sup>40</sup> See Friedman (2017).

<sup>41</sup> Recall we can remain neutral on what exactly the epistemic standard is e.g. truth, knowledge, justification; relative to the standard of truth, or knowledge or justification, I ought to believe 'I am wearing blue socks or Elvis Presley was an alien'.

(I used available inferences and ‘ought<sub>A</sub>-actually’. For completeness, note that the ‘rationalized’ version is also true, where we remove the restrictions to available inferences:

*Rational Epistemic-Requirement (RER)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S epistemically-ought<sub>A</sub>-rationally to infer Q

If *NER* and *RER* seem implausible, it might be because for humans there is always some cost in time, energy or computing power to making an inference. But imagine a creature for whom there was no cost e.g. angels with infinite computing power. If they are at all interested in truth, knowledge or justification, then they would instantaneously make all the inferences from their beliefs. And we could explain the rationality of their doing so in terms of the epistemic ought. Although humans are not like this, I think it is natural to invoke such ideals.

For further support, Christensen (2004, pp. 165–166) gives the following example:

Efficiency seems to enter into the evaluation of car designs in a fairly simple way: the more efficient a car is, the better. Now suppose someone objected to this characterization as follows: “Your evaluative scheme imposes an unrealistic standard. Are you trying to tell me that the Toyota Prius hybrid, at 49 mpg, is an “inefficient” car? On your view, the very best car would use no energy at all! But this is technologically impossible...the very laws of physics forbid it!”

Christensen points out that this objection fails to undermine our ideal of efficiency, concluding that there is room for unattainable ideals even in the most pragmatic endeavours, and that we can recognize the normative force of ideals whose realization is far beyond human capacities.<sup>42</sup>

Moving on, the sense of ‘ought’ in which it is *not* the case that you ought to infer all the trivial logical consequences of your beliefs is the practical sense:

*False Practical-requirement (FPR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-available-inferences, then S *practically*-ought<sub>A</sub>-actually to infer Q<sup>43</sup>


As long as S has limited cognitive capacity and reasons to do things other than believe truths (e.g. to eat, to reproduce) as all known agents do, then FPR will have counterexamples. The problems of clutter avoidance involve such counter-examples. ,<sup>44,45</sup>

<sup>42</sup> This sentence borrows from Christensen’s phrasing on p. 165 and p. 167.

<sup>43</sup> This might be true using practically-ought-rationally, if ideally rational agents can make every inference instantly with no cost. But my point is that we can find a reading which vindicates the intuitions behind the no-clutter objection.

<sup>44</sup> McHugh and Way (2018) makes a similar point in terms of attributive uses of ‘good’: ‘Any ‘intuition’ we have that such beliefs are worthless is likely to concern some form of goodness [all things considered] rather than goodness qua [belief i.e. epistemic goodness]’ p. 24.

<sup>45</sup> It remains an open question what you practically-ought to believe—which inferences practically-should you make? It will depend on your views about the practical standard and the contingent facts about reaching



**Fig. 3** Worlds re-ordered by the practical goal

We can see in Fig. 3 that relative to the practical standard, the best world could be w1, where only obvious inferences are made and the agent can spend their resources doing something else.

For practical ought claims to be true, the agent must have practical goals such that it is worth making the inferences. So the true norm is something like:

*Non-rational Practical-requirement (NPR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-available-inferences, and it is worth S making the inferences then S practically-ought<sub>A</sub>-actually to infer Q

(For completeness, note that the ‘rationalized’ version is also true, where we remove the restrictions to available inferences:

*Rational Practical-requirement (RPR)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and it is worth S making the inference then S practically-ought<sub>A</sub>-rationally to infer Q)

I suggest that NER, RER, NPR and RPR express the link between logic and reasoning.

## 8 Ideals, guidance, appraisal

Norms can be used for i) expressing ideals, ii) for guidance and iii) for making appraisals.<sup>46</sup> But different norms seem to be required for each role. I want to show that our intuitions about the divergence of norms for ideals, guidance and appraisal can be accounted for by the two parameters.

Footnote 45 continued

this standard in your situation (e.g. fulfil desires for Humeans, maximize value for non-Humeans etc.) We should not expect general answers to these questions.

<sup>46</sup> This is the focus of Steinberger (Steinberger 2019a). Compare Kiesewetter (2017, p. 13) who distinguishes first-person guidance, second-person *advice* and third-person *criticism*. I suggest that talk of first-person, second-person and third-person perspectives are ways of making salient different values for the possibilities parameter.

## 8.1 Ideals

Let's start with ideals, which are closely related to standards. Think of the ideal norm as expressing the best way of achieving a given standard, making no allowance for any limitations of an agent or other standards the agent might have. We'll focus on the epistemic standard e.g. believing all and only truths, so the relevant ought is epistemic-ought. As any limitations of the agent are irrelevant for the ideal norm, we need ought-rationally. Putting this together the ideal norm is:

*Rational Epistemic-Requirement (RER)*<sup>47</sup>

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, then S epistemically-ought<sub>A</sub>-rationally to infer Q

For precedent, compare the utilitarian thesis that an act is right if and only if it maximizes happiness. Faced with the objection that this norm fails to provide guidance, utilitarians can maintain that their principle expresses the ideal norm relative to the moral standard, even if we cannot always follow it.<sup>48</sup>

There is a controversy worth mentioning before we go further. What is ideal reasoning for an agent who falsely believes that the relevant rule is invalid? For example, suppose S has been told by a confused teaching assistant that modus ponens is invalid. This is misleading higher order evidence. Should agents reason in line with their false beliefs? Some say no, that misleading higher order evidence should be ignored in first-order reasoning (level-splitters and right reasons theorists<sup>49</sup>). Others say yes (conciliationists). There is an analogous debate in ethics. Some hold that those with misleading higher order evidence about the ethical rules should (morally) ignore that misleading higher order evidence.<sup>50</sup><sup>51</sup>

I have my own views on this controversy (Bradley 2019), but this framework allows us to remain neutral. At the end of Sect. 4 I argued that we can bracket the rest of the agent's epistemic states, and in particular the question of whether P is justified, and focus on the inference from P to Q. Similarly, we can bracket any of the agent's epistemic states that might defeat the inference i.e. make the inference from P to Q incorrect. The advice of a confused teacher would thereby be bracketed. Thus, I leave open the question of how, if at all, contextualism interacts with the debate about higher level evidence.

<sup>47</sup> One possible counter-example is an epistemic Pascal's wager, where the agent will be rewarded with lots of true beliefs if they fail to infer Q (Berker 2013). I set these cases aside here and address them in (ms).

<sup>48</sup> See Bales (1971), Railton (1984), Jackson (1991).

<sup>49</sup> See Horowitz (2014).

<sup>50</sup> Compare Harman (2011).

<sup>51</sup> Whatever line you take here will filter down to norms of guidance and appraisal. For example, if the ideal is that misleading higher order evidence is to be taken into account in first-order reasoning, then it is natural to judge agents who do not take it into account epistemically criticism-worthy. And if the ideal is that misleading higher order evidence is not to be taken into account in first-order reasoning, then it is natural to judge agents who do take it into account epistemically criticism-worthy. And there is room for mixed verdicts. One might hold that the ideal says that misleading higher order evidence is not to be taken into account in first-order reasoning, but agents who do are excused i.e. not criticism-worthy.

Whatever the ideal is, we can now ask how the norms of guidance and appraisal diverge from it.

## 8.2 Guidance

We expect that agents can be guided by norms, but ideal norms cannot always serve as norms of guidance. For example, a norm might say ‘if the exam asks for the capital of Portugal, then write ‘Lisbon’’. This expresses the ideal, but cannot guide an agent who doesn’t know it. (Perhaps better: Doesn’t believe it). In ethics, utilitarians accept that their theory needs to say something about guidance, and they offer norms that can be used to guide e.g. maximize expected utility. In both cases, the natural solution is to hold that norms which can guide agents are restricted to refer only to beliefs and abilities the agent has.

Let’s again focus on the epistemic standard. S can only be guided by inferences available to S, so I suggest that the guidance norm is:

### *Non-rational Epistemic-Requirement (NER)*

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-available-inferences, then S epistemically-ought<sub>A</sub>-actually to infer Q.

## 8.3 Appraisal

What is it to be blameworthy for violating an epistemic norm?<sup>52</sup> Here is a useful principle adapted from Kauppinen (2018):

### *Epistemic Blameworthiness*

S is blameworthy for violating an epistemic norm if and only if it appropriate, other things being equal, to hold the subject accountable by reducing epistemic trust, insofar as she lacks an excuse.

I’m going to assume that *Epistemic Blameworthiness* is roughly correct. My aim in this section is to map our intuitions about what counts as an excuse onto the contextualist framework.

<sup>52</sup> See Steinberger (2019a, p. 17) for more on the epistemic sense of appraisal e.g. ‘Appraisals are central to our normative practices. As normatively regulating and regulated beings, we constantly hold each other accountable for what we do, intend or believe. This is true in the theoretical as much as in the practical arena. In particular, we appraise our epistemic peers in order to assess their character as inquirers. A conscientious reasoner is generally a trustworthy source of testimony. In arriving at such assessments, we frequently care about more than merely an evaluation of their beliefs as correct or mistaken. After all, knowing that my peer arrived at a mistaken belief not out of irresponsible doxastic conduct, but out of sheer bad luck and despite having discharged her epistemic obligations, is valuable information in forming an opinion about her epistemic character.’ See also Smith (2005), McCormick (2011), McHugh (2013) and Brown (2020).



Distinguish two types of excuse for failing to make a valid inference.<sup>53</sup> Agents can be excused by being *unable* to make the inference, or by having *no sufficiently good reason* to make the inference. These excuses correspond to the two parameters. Let's go through them. (I leave open that there might be other types of excuses. I give sufficiency conditions for excuses. Blameworthiness requires no excuses, so I give necessary conditions on blameworthiness.<sup>54</sup>)

First, S might be *unable*<sup>55</sup> to make the inference because it is too complicated, and could thereby be excused.<sup>56</sup> In the contextualist framework, they still infer as they ought-actually to. So if the inference they fail to make is not one they ought-actually to make then they have an excuse.

A complication is that we might reduce our epistemic trust in an agent precisely because they are unable to make the valid inference. For example, it might be an inference that we can make, and which we expect others to make, so S's inability to make that inference reduces our epistemic trust in S. The effect of this complication is to expand the live possible worlds to an intermediate level. For example, consider an agent who can only infer Q1, producing w1. Although they cannot infer Q2 and thereby produce w2, we expect them to be able to, while we do not expect them to infer Q3 and produce w3. So the best live world is w2, and S ought to produce it. Call this middling sense 'ought-competently'. S might fail to infer Q3, but if S infers as they ought-competently then they have an excuse (Fig. 4).<sup>57</sup>

Second, S might have *no sufficiently good reason* to make the inference (because S has non-epistemic goals), and would thereby be excused. In the contextualist framework, they still infer as they practically-ought to. Once non-epistemic goals are added, the ordering of worlds can change, and the best world might be one in which the agent does not make the inference e.g. when the inference is trivial. In Fig. 5, S would be excused for failing to arrive at w2, as the best world is w1; in failing to arrive at w2 or w3, S infers as she practically-ought to. So if the inference they fail to make is not one they practically-ought to make then they have an excuse.

<sup>53</sup> Can there be an epistemic excuse for making an invalid inference? This raises epistemic Pascal's Wager issues which I set aside here.

<sup>54</sup> A plausible further necessary condition is that the speaker endorses the standard (Worsnip 2019a).

<sup>55</sup> Mapping to the debate about epistemic conditions on responsibility, this might be equivalent to the 'control condition' (Rudy-Hiller 2018).

<sup>56</sup> Perhaps some inferences could always be made. If all agents are able to infer according to modus ponens then we get what Broome (2000) calls 'strict liability' for simple logical relations: "The relation between believing p and believing q [a logical consequence of p] is strict. If you believe p but not q, you are definitely not entirely as you ought to be" (85). With a more complicated inference which an agent cannot make, perhaps they *ought to have been able* to make the inference e.g. perhaps they should have taken a logic course. (This is the 'tracing condition'; see Vargas 2005). Then they might be blameworthy for not taking the logic course, and they might be blameworthy for the downstream consequences e.g. being unable to make the inference. It seems to me that consideration of such facts creates a context in which it is true to say that they *could* have made the inference. .

<sup>57</sup> Perhaps a similar point applies to forgetting. If S forgets p then they are unable to believe/infer that p. We would reduce our trust in them only if we expect them to remember p.




Fig. 4 Live worlds for the competent

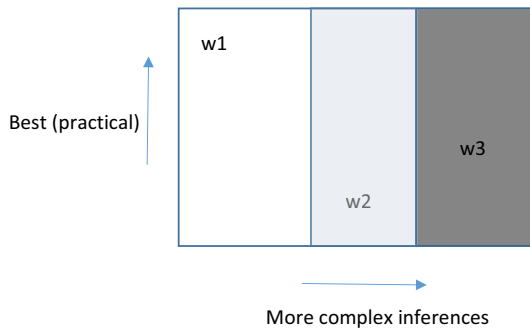


Fig. 5 Worlds re-ordered by the practical goal

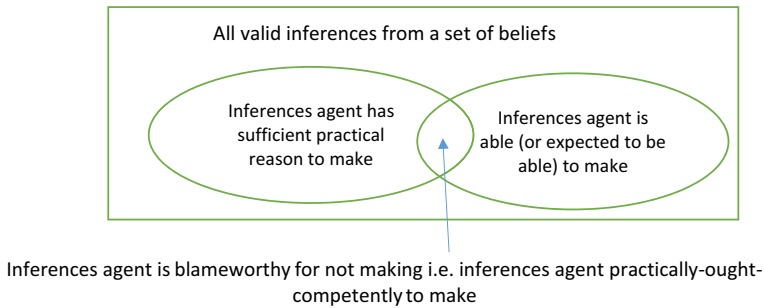
Putting these together, if the inference they fail to make is either not one they practically-ought to make or not one they ought-competently make, then they have an excuse. Contrapositively, if agents are blameworthy for failing to make a valid inference then they fail to infer as they practically-ought-competently. So the norm of blame for reasoning is:

*Non-rational Practical-Requirement+* (NPR+)<sup>58</sup>

For all agents S, and propositions P and Q:

If P entails Q, and S believes P, and P supports Q via-S-competent-inferences, and it is worth S making the inferences, then S practically-ought<sub>A</sub>-competently to infer Q

<sup>58</sup> I've added + because 'competently' has been introduced instead of 'actually'.



**Fig. 6** A Way to be Epistemically Blameworthy

Let's try a case:

### *Melted Ice-cream*

Alessandra has gone to pick up her children at their elementary school. It is hot, but she leaves the ice-cream she has brought for her children in the car. Although able to infer that the ice-cream will melt, she does not do so. By the time they return the ice-cream has melted.<sup>59</sup>

Intuitively, Alessandra is epistemically blameworthy. We would reduce our epistemic trust in Alessandra if we learnt that she failed to realize that the ice-cream would melt. Our framework delivers this verdict if the inference to the belief that the ice-cream would melt is one she is both able to make and has sufficient practical reason to make. And indeed both conditions are satisfied. Alessandra has enough inferential competence to be able to work out that the ice-cream would melt, and has sufficient practical interest in the ice-cream not melting.<sup>60</sup> She practically-ought-competently to have inferred that the ice-cream would melt, but she does not, so is epistemically blameworthy (Fig. 6).

Alessandra is excused if we make either of two modifications to the story. If we modify the story to one in which her full attention on something other than the ice-cream is a matter of life and death then Alessandra is not epistemically blameworthy. For example, suppose she is a doctor and as she parks she sees that there has been an accident and only her full attention for several hours will save the life of a child. In such a context, a melting ice-cream is trivial in the same sense that it is trivial to infer that I am wearing blue socks or Elvis is alive. She does not have practical reason to make the inference, so is not blameworthy.<sup>61</sup>

Alessandra is also excused if the inference to the belief that the ice-cream would melt is not one she is able to make, nor one we would expect her to make. This requires a bit more imagination, but we could imagine that it is a typically cold day in the Arctic

<sup>59</sup> Adapted from Sher (2009, p. 24). I've changed the dog to an ice-cream to take morality out of it.

<sup>60</sup> Compare Lillehammer (2019).

<sup>61</sup> See Schroeder (2012) for related examples where there is good reason not to deliberate. Similar issues arise for the question of when it is rational to reconsider a belief (see Paul 2015). Fragmented agents (Stalnaker 1984) can perhaps be thought of as agents who do not make the inference which would connect the fragments.

Circle where the ice-cream would normally not melt, but the car is parked in a place where heat from concave neighbouring buildings is focussed. Alessandra knows the contingent facts, but does not have the mathematical abilities necessary to work out that the ice-cream would melt. She ought-rationally to make the inference, but we would not expect her to be able to make the inference, so she is not blameworthy.

## 9 Conclusion

I have argued that many controversies about the norms of reasoning can be resolved by an independently motivated contextualist semantics for ‘ought’. The problems of belief revision and the preface paradox can be solved by relativizing to a set of propositions, the problem of excessive demands can be solved by relativizing to a set of available inferences, and the problem of clutter avoidance can be solved by relativizing to a standard. These parameters can also illuminate questions about which norms are relevant to ideals, guidance, and blame.<sup>62</sup>

**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/>.

## References

- Allo, P. (2016). Logic, reasoning and revision. *Theoria*, 82, 3–31. <https://doi.org/10.1111/theo.12071>.
- Arpaly, N. (2002). *Unprincipled virtue: An inquiry into moral agency*. Oxford: Oxford University Press.
- Bales, R. E. (1971). Act-utilitarianism: account of right-making characteristics or decision-making procedure? *American Philosophical Quarterly*, 8(3), 257–265.
- Berker, S. (2018). A combinatorial argument against practical reasons for belief. *Analytic Philosophy*, 59(4), 427–470.
- Björnsson, G., & Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121(1), 7–36.
- Boghossian, P. A. (2003). The normativity of content. *Philosophical Issues*, 13, 31–45.
- BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press.
- Bradley, D. (2019). Are there indefeasible epistemic rules? *Philosopher’s Imprint*, 19(3), 1–19.
- Broome, J. (1999). Normative requirements. *Ratio*, 12(4), 398–419.
- Brogaard, B. (2008). Moral contextualism and moral relativism. *Philosophical Quarterly*, 58(232), 385–409.
- Broome, J. (2000). Normative requirements. In J. Dancy (Ed.), *Normativity*. Oxford: Blackwell.
- Broome, J. (2013). *Rationality through reasoning*. New York: Wiley.
- Brown, J. (2020). Epistemically blameworthy belief. *Philosophical Studies*, 177, 3595–3614.
- Cariani, F., Kaufmann, M., & Kaufmann, S. (2013). Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy*, 36(3), 225–259.
- Carr, J. (2015). Subjective ought. *Ergo*, 2(27), 678–710.
- Charlow, N. (2013). What we know and what to do. *Synthese*, 190(12), 2291–2323.

<sup>62</sup> Thanks to Kristoffer Ahlstrom-Vij, Stephen Finlay, Jess Isserow, Anti Kauppinen, Mona Simion, Pekka Vayrynen, Daniel Whiting and several referees for comments and discussion of this paper. This research was partly funded by the British Academy Mid-Career Fellowship MD19\190046.

- Chignell, A. (2018) The ethics of belief. *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (Ed.), <https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>.
- Chisholm, R. M., et al. (1989). *Theory of knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Chrisman, M. (2015). *The meaning of 'Ought': Beyond descriptivism and expressivism in metaethics*. Oxford: Oxford University Press.
- Christensen, D. (1994). Conservatism in epistemology. *Noûs*, 28(1), 69–89.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. Oxford University Press.
- Cowie, C. (2014). In defence of instrumentalism about epistemic normativity. *Synthese*, 191(16), 4003–4017.
- DePaul, M. R. (2004). Truth consequentialism, withholding and proportioning belief to the evidence. *Philosophical Issues*, 14(1), 91–112.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- Douven, I. (2010). The pragmatics of belief. *Journal of Pragmatics*, 42(1), 35–47.
- Dowell, J. (2012). Contextualist solutions to three puzzles about practical conditionals. *Oxford Studies in Metaethics*, 7, 271–303.
- Dowell, J. (2013). Flexible contextualism about deontic modals: A puzzle about information-sensitivity. *Inquiry*, 56(2-3), 149–178.
- Dreier, J. (1990). Internalism and speaker relativism. *Ethics*, 101, 6–26.
- Engel, P. (2007). Belief and normativity. *Disputatio*, 23, 179–203.
- Fassio, D. (2019) 'The aim of belief'. *The Internet Encyclopedia of Philosophy*, <https://www.iep.utm.edu/beli-aim/>
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60, 667–695.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15–34.
- Field, H. (2009). I—What is the normative role of logic? In *Aristotelian society supplementary volume* (Vol. 83, No. 1, pp. 251–268). Oxford: Blackwell Publishing Ltd.
- Field, H. (2015). What is logical validity? In C. Caret & O. Hjortland (Eds.), *Foundations of logical consequence*. Oxford: Oxford University Press.
- Finlay, S. (2014). *Confusion of Tongues: A theory of normative language*. Oxford: Oxford University Press.
- Foot, P. (1972). Morality as a system of hypothetical imperatives. *Philosophical Review*, 81(3), 305–316.
- Frankena, W. K. (1958) Obligation and motivation in recent moral philosophy. In A. I. Melden (Eds.), *Essays in moral philosophy* (Seattle, pp. 40–81).
- Frege, G. (1893/1903/2009). *Grundgesetze der Arithmetik*. Paderborn: Mentis
- Friedman, J. (2017). Junk beliefs and interest-driven epistemology. *Philosophy and Phenomenological Research*, 97(3), 568–583.
- Friedman, J. (forthcoming) The epistemic and the zetetic. *The Philosophical Review*.
- Geach, P. (1956). Good and evil. *Analysis*, 17, 33–42.
- Geach, P. T. (1982). Whatever happened to deontic logic? *Philosophia*, 11(1–2), 1–12.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge, MA: Harvard University Press.
- Gibbons, J. (2013). *The norm of belief*. Oxford University Press.
- Goldman, A. I. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- Good, I. J. (1966). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4), 319–321.
- Greco, (2015). Epistemological open questions. *Australasian Journal of Philosophy*, 93(3), 509–523.
- Harman, G. (1973). Review of Roger Wertheimer, the significance of sense: Meaning, modality, and morality. *The Philosophical Review*, 82(2), 235–239.
- Harman, G. (1975). Moral relativism defended. *The Philosophical Review*, 84, 3–22.
- Harman, G. (1984). Logic and reasoning. *Synthese*, 60, 107–127.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
- Harman, G. (1996). Moral relativism. In G. Harman & J. Thomson (Eds.), *Moral relativism and moral objectivity*. Oxford: Blackwell.
- Harman, E. (2011). Does moral ignorance exculpate? *Ratio*, 24(4), 443–468.
- Hedden, B. (2015). *Reasons without persons: Rationality, identity, and time*. Oxford: Oxford University Press.
- Horowitz, S. (2014). Epistemic Akrasia. *Noûs*, 48(4), 718–744.
- Hughes, N. (2017). Dilemmic epistemology. *Synthese*, 1–32.
- Jackson, E. (2018). Belief, credence, and evidence. *Synthese*, 1–20.

- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461–82.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
- Kauppinen, A. (2018). Epistemic norms and epistemic accountability. *Philosophers' Imprint*, 18.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research*, 66(3), 612–640.
- Khoo, J., & Knobe, J. (2018). Moral disagreement and moral semantics. *Noûs*, 52(1), 109–143.
- Kiesewetter, B. (2017). *The normativity of rationality*. Oxford University Press.
- Knobe, J., & Szabó, Z. G. (2013). Modals with a taste of the deontic. *Semantics and Pragmatics*, 6(1), 1–42.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455), 509–563.
- Kolodny, N., & MacFarlane, J. (2010). Ifs and Oughts. *Journal of Philosophy*, 107(3), 115–143.
- Kopec, M. (2018). A pluralistic account of epistemic rationality. *Synthese*, 195(8), 3571–3596.
- Kratzer, A. (1981). The notional category of modality. In H.-J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts: New approaches in world semantics* (pp. 38–74). Berlin: de Gruyter.
- Kratzer, A. (1991). Modality. In A. von Stechow & D. Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research* (pp. 639–650). Berlin: de Gruyter.
- Kratzer, A. (2012). *Modals and conditionals*. Oxford: Oxford University Press.
- Kyburg, H. E. (1961). *Probability and the logic of rational belief*. Middletown, CT: Wesleyan University Press.
- Lewis, D. (1980). A Subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in inductive logic and probability* (Vol. 2, pp. 263–293). Berkeley, CA: University of California Press.
- Lewis, D. (1986). *Philosophical papers* (Vol. II). Oxford: Oxford University Press.
- Lillehammer, H. (2019) Review of 'The Importance of Being Rational' <https://ndpr.nd.edu/news/the-importance-of-being-rational/>
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- MacFarlane, J. (2004) (ms). In what sense (if any) is logic normative for thought? (unpublished).
- Makinson, D. C. (1965). The paradox of the preface. *Analysis*, 25(6), 205–207.
- McHugh, C. (2012). The truth norm of belief. *Pacific Philosophical Quarterly*, 93, 8–30.
- McHugh, C. (2013). Epistemic responsibility and doxastic agency. *Philosophical Issues*, 23(1), 132–157.
- McHugh, C., & Way, J. (2018). What is good reasoning? *Philosophy and Phenomenological Research*, 96(1), 153–174.
- McCormick, M. (2011). Taking control of belief. *Philosophical Explorations*, 14(2), 169–183.
- Moore, G. E. (1903). *Principia ethica*. Cambridge: Cambridge University Press.
- Nebel, J. M. (2019). Normative reasons as reasons why we ought. *Mind*, 128(510), 459–484.
- Paul, S. K. (2015). Doxastic self-control. *American Philosophical Quarterly*, 52(2), 145–158.
- Perry, J. (1979). The problem of the essential indexical. *Noûs*, 3–21.
- Prawitz, D. (2005). Logical consequence from a constructivist view. In *The Oxford handbook of philosophy of mathematics and logic*.
- Quine, W. V. O. (1970). Methodological reflections on current linguistic theory. *Synthese*, 21(3/4), 386–398.
- Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy and Public Affairs*, 13(2), 134–171.
- Rinard, S. (2019). Equal treatment for belief. *Philosophical Studies*, 176(7), 1923–1950.
- Rudy-Hiller, F. (2018). The epistemic condition for moral responsibility. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2018 Edition), <https://plato.stanford.edu/archives/fall2018/entries/moral-responsibility-epistemic/>.
- Ryan, S. (1999). The logic of rationality. *Philosophia*, 27, 287–299.
- Sainsbury, M. (2002). What logic should we think with? In A. O'Hear (Ed.), *Royal institute of philosophy lecture series*. Cambridge: Cambridge University Press.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18, 337–364.
- Schroeder, M. (2011). Ought, agents, and actions. *Philosophical Review*, 120(1), 1–41.
- Schroeder, M. (2012). The ubiquity of state-given reasons. *Ethics*, 122(3), 457–488.
- Schwarz, W. (2020). Ability and possibility. *Philosophers' Imprint*, 20, 6.
- Shafer-Landau, R. (2005). Error theory and the possibility of normative ethics. *Philosophical Issues*, 15(1), 107–120.

- Shah, N. (2003). How truth governs belief. *The Philosophical Review*, 112, 447–482.
- Shah, N., & Velleman, J. D. (2005). Doxastic deliberation. *The Philosophical Review*, 114, 497–534.
- Sher, G. (2009). *Who knew: Responsibility without awareness*. New York: Oxford University Press.
- Silk, A. (2014). Why 'Ought' detaches: Or, why you ought to get with my friends (if you want to be my lover). *Philosophers*, 14(7), 1–16.
- Silk, A. (2017). Normative language in context. *Oxford Studies in Metaethics*, 12, 206–239.
- Simion, M., Kelp, C., & Ghijssen, H. (2016). Norms of belief. *Philosophical Issues*, 26(1), 374–392.
- Singer, D. J. & Aronowitz, S. (forthcoming). What epistemic reasons are for: Against the belief-sandwich distinction. In B. Dunaway & D. Plunkett (Eds.), *Meaning, decision, and norms: Themes from the work of Allan Gibbard*. Available on PhilArchive: <https://philarchive.org/archive/SINWER>
- Sinnott-Armstrong, W. (2014). Do psychopaths refute internalism? In Schramme, T. (Ed.). *Being amoral: Psychopathy and moral incapacity* (pp. 187–208). MIT Press.
- Smith, A. M. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics*, 115(2), 236–271.
- Stalnaker, R. (1984). *Inquiry*. Cambridge University Press.
- Steinberger, F. (2016). Explosion and the normativity of logic. *Mind*, 125(498), 385–419.
- Steinberger, F. (2019a). Three ways in which logic might be normative. *The Journal of Philosophy*, 116(1), 5–31.
- Steinberger, F. (2019b). Logical pluralism and logical normativity. *Philosopher's Imprint*, 19(12), 1–19.
- Steinberger, F. (2019c). Consequence and normative guidance. *Philosophy and Phenomenological Research*, 98(2), 306–328.
- Sturgeon, S. (1998). Humean chance: Five questions for David Lewis. *Erkenntnis*, 49(3), 321–335.
- Titelbaum, M. G. (2015). How to derive a narrow-scope requirement from wide-scope requirements. *Philosophical Studies*, 172(2), 535–542.
- Valaris, M. (2019). Reasoning and deducing. *Mind*, 128(511), 861–885.
- Vargas, M. (2005). The trouble with tracing. *Midwest Studies in Philosophy*, 29(1), 269–290.
- Vetter, B., & Viebahn, E. (2016). How many meanings for 'may'? The case for modal polysemy. *Philosophers' 16*.
- Wallace, R. J. (2001). Normativity, commitment, and instrumental reason. *Philosophers' Imprint*, 1(3), 1–26.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16, 267–297.
- Wedgwood, R. (2006). The meaning of 'ought'. *Oxford Studies in Metaethics*, 1, 127–160.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.
- Wedgwood, R. (2012). Justified inference. *Synthese*, 189(2), 1–23.
- Wedgwood, R. (2016). Objective and subjective 'ought'. *Deontic Modality*, 143–168.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19(1), 445–459.
- Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society*, 39, 103–138.
- Williamson, T. (1998). Conditionalizing on knowledge. *British Journal for the Philosophy of Science*, 49(1), 89–121.
- Woods, J., & Maguire, B. (forthcoming). The Game of Belief *Philosophical Review*
- Worsnip, W. (2019a). 'Ought'-contextualism beyond the parochial. *Philosophical Studies*, 176(11), 3099–3119.
- Worsnip, W. (2019b). Isolating correct reasoning. In M. B. Jackson & B. B. Jackson (Eds.), *Reasoning: New Essays on theoretical and practical thinking*. Oxford: Oxford University Press.
- Worsnip, A. (2020). Resisting relativistic contextualism: On finlay's confusion of tongues. *Analysis*, 80(1), 122–131.