

## OUGHT-IMPLIES-CAN IN CONTEXT

DARREN BRADLEY

*University of Leeds*

If you ought to do something, does it follow that you are able to do it? The Kantian thesis that ought-implies-can seems intuitive and is widely accepted. Nevertheless, there are several powerful purported counterexamples. In this paper I will apply an independently motivated contextualism about 'ought' to make sense of the intuitions on both sides of the argument. Contextualism explains why ought-implies-can seems compelling despite being false in many contexts. The result will be that philosophers cannot in general appeal to ought-implies-can in their arguments.

### 1. Introduction

If you ought to do something, does it follow that you are able to do it? The Kantian thesis that ought-implies-can seems intuitive and is widely accepted.<sup>1</sup> It is neatly expressed in the Latin motto: *Ad impossibilia nemo tenetur*: No-one is bound to do the impossible. Nevertheless, there are several powerful purported counterexamples. In this paper I will apply an independently motivated contextualism about 'ought' to the debate. The idea is that, as 'ought' has different meanings depending on the context in which it is uttered, there are contexts where 'ought-implies-can' is true and contexts where it is false. Specifically, we can distinguish an *ideal* sense of 'ought', a *deliberative* (or guidance) sense of 'ought', and a *hypological* (involving praise or blame) sense of 'ought'. I will argue that ought-implies-can is true in the deliberative sense, false in the ideal sense, and false in the hypological sense unless the agent is blameworthy for the inability.

---

1. See Moore (1922) and Stocker (1971) for arguments that ought-implies-can is intuitively correct. Recent defenders include Fischer (2003), Streumer (2003), Howard-Snyder (2006), Vranas (2007), Copp (2008) and Chuard & Southwood (2009). Buckwalter (2020: 23) writes: "A widely accepted view in Western philosophy is that moral obligations entail the ability to fulfill them, often glossed in the slogan that ought implies can."

**Contact:** Darren Bradley <bradleydarren@gmail.com>

It is surprising that despite the large literature on ought-implies-can and the large literature on the context-sensitivity of normative terms there has been almost no discussion of how they might interact. The only paper I know of which develops the connection is Berg (2018). However, Berg doesn't explain how contextualism interacts with the arguments for and against ought-implies-can. I will defend a theory which makes the connections explicit, and diverge from Berg on the treatment of blame.

In §2 I review the main arguments for ought-implies-can, and in §3 the main arguments against. In §4 I explain the contextualist framework and three hypotheses for how the context determines the meaning of 'ought'. In §5 I apply the contextualist framework to the arguments of §§ 2 and 3, and show how our intuitions are predicted by the three hypotheses. In § 6 I discuss the problem of disagreement for contextualism, in §7 some dialectical consequences, and in §8 I compare my account to alternative ways of weakening ought-implies-can.

Let me make five quick clarifications. First, I will follow the literature in assuming that 'ought to' and 'obligated to' are synonymous.

Second, I will use contextualism for concreteness, but there are alternatives to contextualism regarding the linguistic implementation of the semantic variability.<sup>2</sup> Arguments that contextualism is the correct linguistic theory are to be found in the linguistics literature and I have nothing to add to them. The power of contextualism lies in its ability to explain our competing intuitions. I will argue that contextualism explains our intuitions, and I take this to be a powerful indirect argument in favour of it.

Third, I distinguish the view that 'ought' is context-sensitive from the view that 'can' is context-sensitive, and focus on the former. While I agree that 'can' is context-sensitive, and that this is relevant to whether ought-implies-can,<sup>3</sup> I think that the context-sensitivity of 'ought' is crucial and under-discussed, hence the focus of this paper.

Fourth, many non-contextualists agree that ought-implies-can is ambiguous, and comes in different degrees of logical strength.<sup>4</sup> The logically stronger readings can be used to explain cases where ought-implies-can seems false and logically weaker readings can be used to explain cases where ought-implies-can seems true. On my view, the *context* explains these intuitions, so my use of ought-implies-can can be understood in the logically stronger sense. I hold that in contexts where it is true that ought-implies-can, a logically strong reading

---

2. See Vetter & Viebhan (2016).

3. See Norcross (2020: 133–4) for a clear application to ought-implies-can. For related discussions see Schwarz (2020), Suikkanen (2020), Bradley (2021), and Talbot (2021).

4. See King (2019) and Bassford (2022) for numerous choice-points.

is still true (and in contexts where it is false that ought-implies-can, a logically strong reading will be all the more obviously false).<sup>5</sup>

Finally, many of my uses of ‘ought’ are mentions rather than uses, and really belong in inverted commas. For the sake of readability, I will often omit the inverted comments (see Lewis 1996: 566–7 for the same strategy).

## 2. For Ought-Implies-Can

In this section I will briefly explain two arguments for ought-implies-can.<sup>6</sup>

### 2.1 *The Argument from Deliberation*

One argument for ought-implies-can is that it explains why one cannot rationally deliberate about impossible actions. Hare writes:<sup>7</sup>

Suppose... that I am being driven on the French shore by a gale and that it is obvious that whether I shall land in France is out of my control; then, if I ask ‘Shall I land in France?’ this cannot be understood as a practical question... It is, in fact, the impossibility of deliberating, or wondering, whether to do a thing which rules out asking whether one ought to do it. (1963: 60–1)

The impossibility of (rationally) deliberating about something which is impossible is explained by ought-implies-can. If you cannot do something then it is not the case that you ought to do it, so you should not deliberate about whether to do it.

Pictorially, we can imagine a set of epistemically possible worlds which some agent *can* produce, and a set of epistemically possible worlds which the agent *ought* to produce. Advocates of ought-implies-can hold that the worlds the agent *ought* to make true are a sub-set of the worlds the agent *can* make true.

---

5. To give one example (King’s [2019] second choice point), I will assume that ‘implies’ refers to strict implication rather than some weaker implication such as conversational implicature. Where neither choice is logically stronger than the other, I intend what I say to apply to both.

6. There are surprisingly few. Talbot (2016: 378) writes: ‘In many cases, however, those who endorse ought-implies-can do not really argue for the principle, but focus on responding to objections (e.g., Zimmerman 1996; Vranas 2007).’

7. See Williams (1979), Howard-Snyder (2006), Streumer (2007), Vranas (2007), and Henne et al. (2019) for related arguments.




Figure 1: Ought implies can.

## 2.2 *The Argument from Excuses*

Another argument for ought-implies-can is based on the intuition that inability provides an excuse. A purported obligation seems to be excused if it turns out that it cannot be performed. And if you are excused, it is not the case that you ought to do it. Copp writes:

If marauding sharks would attack Aurelia and prevent her from successfully rescuing a drowning man, then she cannot rescue him, and this seems enough to show that she is not morally required to rescue him. Ought implies can seems to explain our intuitions. (1997: 446)

Call this case *Sharks*. Such cases give strong support to ought-implies-can.

There are other arguments for ought-implies-can of course. One further argument is based on linguistic practice (see Sinnott-Armstrong 1984). Several others are discussed in Buckwalter (2020). But the two arguments above will be sufficient to demonstrate how the contextualist approach works.

## 3. *Against Ought-Implies-Can*

In this section I briefly explain four arguments against ought-implies-can.

### 3.1 *The Objection from Making it Impossible for Oneself to do Something*

The first objection is a direct response to the argument from excuses. Recall the argument from excuses said that inability to perform an action provides an excuse. The objection is that this makes excuses too easy to come by. Specifi-

cally, the argument from excuses suggests that agents are excused by *making themselves* unable to perform the action. Peter Vranas, although a proponent of ought-implies-can, opens his paper with the following dialogue:

*Late Paper*

“Good morning, Professor. Unfortunately, I haven’t even started writing a paper yet. I know I promised to turn in a paper today by 9 am, but last night I didn’t feel like writing, so I went to the movies instead. What should I do?”

“You should turn in a paper by 9 am. It’s 8:57, so you’ve got three minutes left.”

“Sorry, Professor, maybe I didn’t make myself clear. I don’t have a paper to turn in.”

“You asked what you should do. Obviously, you should fulfil your obligations. You may not have a paper, but you do have an obligation to turn in a paper by 9 am.”

“But I can’t turn in a paper by 9 am, so I don’t have an obligation to do so: ‘ought’ implies ‘can’.” (2007: 167–168)<sup>8</sup>

We are invited to think that the student cannot get off the hook so easily, that inability does not provide an excuse, and that ought does not imply can. Deniers of ought-implies-can hold that the worlds the agent ought to make true are *not* a sub-set of the worlds the agent can make true.




Figure 2: Ought does not imply can.

8. See Sinnott-Armstrong (1984), Streumer (2007), and Heuer (2010) for similar examples discussion of such cases.

### 3.2 *The Objection from Role-Oughts*

Richard Feldman introduces the concept of role-oughts:

#### *Role-oughts*

There are oughts that result from one's playing a certain role or having a certain position. Teachers ought to explain things clearly. Parents ought to take care of their kids. Cyclists ought to move in various ways. Incompetent teachers, incapable parents, and untrained cyclists may be unable to do what they ought to do. (2000: 676)

It seems that agents with these roles ought to perform certain actions even when unable to do so.

### 3.3 *The Objection from Addictive Behaviours*

Another class of cases in which obligations seem to survive the inability to perform the action involve compulsive behaviours. Vranas writes:<sup>9</sup>

Kleptomaniacs cannot refrain from stealing; pathological sadists cannot refrain from hurting others; alcoholics and nicotine addicts cannot refrain from drinking and smoking respectively; yet all these people have...obligations to refrain from the corresponding behaviors—or so the objection to [ought-implies-can] goes. (2018: 183)

### 3.4 *The Objection from Fixed Feelings*

A related style of counter-example involves feelings rather than actions. It can be natural to say that someone ought to feel a certain way, even if they are unable to do so. Vranas writes:

It is frequently claimed that in some cases an agent cannot make herself feel a certain way—e.g., grateful—right away (because how she feels is not under her instantaneous voluntary control) although it is natural to say that the agent ought to feel that way. (2018: 174)

---

9. See also Ryan (2003), Streumer (2007: 269), Graham (2011: §3), and King (2019).

These four arguments suggest that there can be actions which an agent ought to perform but cannot. Again, I do not intend to provide an exhaustive list of arguments against ought-implies-can. But I think these arguments cover the main points and will suffice to show how the contextualist approach works.<sup>10</sup>

#### 4. Contextualism

In this section I explain the core ideas of normative contextualism and state three hypotheses. It is a familiar thought that whether someone is correctly described as tall depends on the details of the conversation. For example, Michael Jordan, at 1.98m, is tall for an ordinary person, but not tall for a basketball player. So, the truth of ‘Michael Jordan is tall’ depends on the conversational context. It is true given a context in which ordinary people are the being discussed, but false given a context in which basketball players are being discussed.

A popular theory in linguistics is that normative terms like ‘ought’ are context-sensitive in a similar way.<sup>11</sup> Applied to ought-implies-can we get:

For some (not all) contexts C: ‘ought implies can’ is true in C.

The traditional Kratzerian semantics posits two parameters – a set of live possible worlds,<sup>12</sup> and an ordering of how good they are. ‘S ought to  $\Phi$ ’ means that in the best live possible worlds, S  $\Phi$ s.<sup>13</sup> And ought-implies-can means that the worlds the agent *ought* to make true are a sub-set of the worlds the agent *can* make true:

‘Ought implies can’ is true in C for S iff the best live worlds are a sub-set of the worlds S can make true in C.

Putting these together:

---

10. The proponent of context-sensitivity about ‘can’ might defend ought-implies-can by saying that in all four stories agents *can perform the actions in the relevant sense*.

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

12. I will take for granted that propositions can be modelled by possible worlds, and use propositions and sets of possible worlds interchangeably (setting aside complications for possible worlds caused by actions, which might require centred possible worlds, or practitions [Schroeder 2011; Chrisman 2012]).

13. Kratzer (1981).

For some (not all) contexts C: the best live worlds are a sub-set of the worlds S can make true in C.

To give this theory substance we need to say something about the contexts in which ought implies can, and the contexts in which it does not.

#### 4.1 *Two Parameters, Three Hypotheses*

Starting with the first parameter, what determines the live possible worlds? I suggest that part of what determines the live possible worlds is the aim of the conversation.<sup>14</sup> We can distinguish (at least) three aims people might have when using 'ought'<sup>15</sup>:

- i) Expressing ideals.
- ii) Deliberating or offering guidance.
- iii) Ascribing praise or blame (hypological)

To see how (i) and (ii) can come apart, consider the student who, at 8:57am, has not written their paper, which is due at 9am. The Professor, knowing this, might say

“You should hand in your paper by 9am.”

We can understand the Professor's utterance as (i) expressing an ideal. In this sense, the utterance is true and relevant. However, it makes no sense to understand the utterance as (ii) offering guidance, given that the Professor knows that the paper cannot be written in 3 minutes.

To see how (i) and (iii) can come apart, suppose that the student has no paper because they were in a coma for the last month, through no fault of their own. Knowing this, the Professor might reasonably make the same utterance to (i) express an ideal, but not to (iii) ascribe blame.

To see how (ii) and (iii) can come apart, notice that the Professor might reasonably make the utterance to (iii) ascribe blame without intending to (ii) offer guidance, e.g. in *Late Paper*, where the Professor knows the student went to the movies.

---

14. The beliefs and evidence of the participants of the conversation might also be relevant, as well as the actual facts.

15. All three can be found in Steinberger (2019: 7). Relatedly, Bales (1971) distinguishes decision procedures from right-making characteristics; Arpaly (2003: 34) distinguishes a rational agent's manual from an account of rationality; Schroeder (2011 p.1–2) distinguishes deliberative from evaluative oughts; McHugh (2012: 9–10) distinguishes prescriptive norms from evaluative norms; and Schoenfield (2018: 690) distinguishes plans to make from procedures to conform to.



Let's make these aims explicit using 'ideal-ought', 'deliberative-ought', and 'hypological-ought'. I suggest that which 'ought' is used partially determines which possibilities are live. Ideal-ought is associated with the most expansive set, deliberative-ought is associated with the least expansive set, and hypological-ought is associated with an intermediate set. Specifically, I offer three hypotheses which I will defend in the next section by showing how they predict the intuitive verdicts in all the examples:

- a) Ideal-ought does *not* imply can.
- b) Deliberative-ought *does* imply can.
- c) Hypological-ought *does* imply can, *unless* the agent is blameworthy for the<sup>16</sup> inability.<sup>17</sup>

The main options for the second parameter, the ordering of possible worlds, are either *moral*, *practical*, or *epistemic* standards.<sup>18</sup> For example, a moral standard will rank worlds in order of how morally good they are, while a practical standard might rank worlds in order of how good they are relative to the interests of some agent. I mention this for completeness, but I won't make use of any shifts to this parameter. I think the conflicting intuitions about cases can be explained purely by the possible worlds parameter.

## 5. Examples Redux

I'll now show how contextualism plus the three hypotheses predicts all the opposing intuitions.

### 5.1. The Argument from Deliberation

Let's consider the deliberation-based argument for ought-implies-can. Recall the idea is that ought-implies-can explains why we cannot deliberate about impossible actions.

---

16. 'The' is grammatically problematic, as there is no earlier reference to an inability to refer back to. We could re-state the hypothesis as: cannot implies not hypological-ought, unless you are blameworthy for the inability. More fully: For all agents G and actions A, if G cannot A then it is not the case that F hypologically-ought to A, unless G is blameworthy for the inability.

17. Here my framework differs from that of Berg (2018) who posits two types of blame, corresponding to ideal and non-ideal normative claims.

18. For further details, see accounts which relativize sentence to codes of practice (Copp 1995), standards (Silk 2016), ends (Finlay 2014), or motivational attitudes (Harman 1975; Dreier 1990).

But we do not need to invoke an unqualified ought-implies-can. The contextualist can hold that when the question concerns whether an agent should deliberate, we restrict the live possible worlds to include only actions the agent is able to perform.<sup>19</sup> The only restriction needed is that predicted by (b) deliberative-ought implies can.<sup>20</sup>




Figure 3: Deliberative ought implies can.

And we can generate the intuition that ought does not imply can by imagining a context that invokes the ideal-ought. Suppose the plan was to land in France, but in the chaos of the storm, no-one can remember the plan. Even assuming that landing in France is now the only possibility, it still makes sense to wonder “ought we to land in France?” This makes sense if the context is one in which the ideal-ought is being used and we ideally-ought to follow the plan.

## 5.2. *The Argument from Excuses*

Recall that in *Sharks* we are inclined to reason from “Aurelia cannot save the man” to “it is not the case that Aurelia ought to save the man.” I suggest that this intuition is strongest in contexts where the conversation is about either deliberation or blame. Given that Aurelia cannot save the man, (b) correctly predicts that it makes no sense for her to deliberate about whether she should. And given that Aurelia cannot save the man (due to no fault of her own), (c) correctly predicts that it makes no sense for others to blame Aurelia for not saving the man. And a story involving a tragic accident naturally makes us think about whether anyone is blameworthy for the accident, thus creating a context in which the hypo-

19. It is not technically a conversational context, but I don’t see why an actual conversation is necessary. If you think it is, perhaps the agent is having a silent conversation with themselves.

20. See Williams (1979), Schroeder (2011), and Hedden (2012) for discussion of the deliberative ought. For dissent see Hicks (2022).

logical ought is relevant. So we can predict, in this story, that inability provides an excuse.

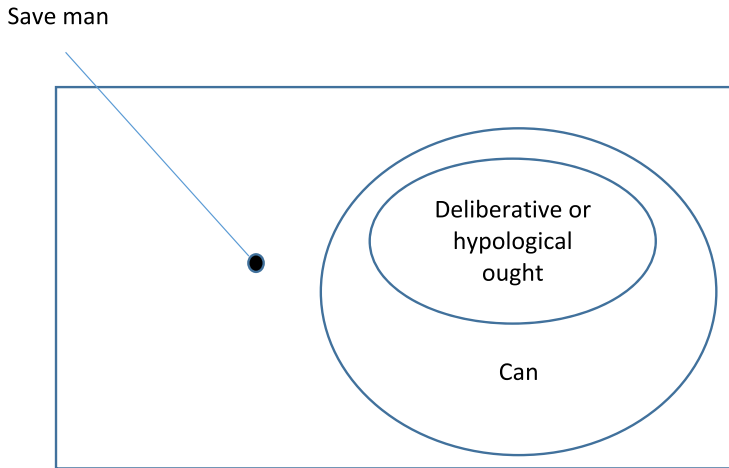


Figure 4: Aurelia cannot save the man and ought implies can.

Perhaps we can generate a different intuition using the same story. Suppose Aurelia, helplessly watching from the beach, says “I ought to save him, but can’t get past the sharks.” We can make sense of this utterance as a true claim about which possibility is best, not making allowances for the practical difficulties. The sentence might be uttered as a regretful comment about the best possibility, which is tragically out of reach. And its truth is predicted by the assumption that Aurelia is expressing the ideal ought, and that (a) ideal-ought does not imply can.

### 5.3. *The Objection from Making it Impossible for Oneself to do Something*

However, there are examples where inability does not seem to provide an excuse, and the contextualist can explain these as involving contexts where the agent is blameworthy for the inability. Recall the student who is unable to hand in the paper because they spent the previous night at the movies. Clearly they are blameworthy for being unable to hand in the paper. The intuition that the student does not have an excuse is explained by the hypothesis that (c) hypothetical-ought does imply can, *unless the agent is blameworthy for the inability*. When they are blameworthy for the inability, as in this case, they can be blameworthy for failing to do something that they cannot do. The result is a set of live possible worlds that includes the action at issue.

Hand in paper on time




Figure 5: Monday 8:57am, wide set of live possible worlds.

And we can generate a different intuition by shifting the context to one that is explicitly about deliberation/guidance. The student might say “I’ve no excuse, but I want to know what I should do now. I can’t write a paper in 3 minutes, so clearly it’s not the case that I ought to hand in the paper by 9...” We can hear this utterance in a way that makes it true. And (b) deliberative-ought-implies-can predicts that this utterance can be true.

Someone might object that the condition *unless the agent is blameworthy for the inability* is ad hoc. But the condition follows from the independently plausible principle that one is responsible for the foreseeable consequences of one’s actions. The principle we need is something like:

If action A foreseeably causes B, which foreseeably causes C, then the agent who A’s is responsible for B and C.<sup>21</sup>

In our example, going to the movies is the action (A), the inability is a mediating cause (B) and the failure to turn in the paper the outcome (C). Perhaps the best way to think about this is that the student is blameworthy for the foreseeable *sequence* of events that their actions cause. *A fortiori*, the student is blameworthy for failing to turn in the paper.

#### 5.4. Role-Oughts

Recall that uses of role-oughts suggest that agents sometimes ought to do things that they are unable to do. The contextualist can explain this by positing that

---

<sup>21</sup>. See Zimmerman (1997), Rosen (2004), Miller (2017) for discussion of related principles. There is controversy over whether logically stronger versions of the principle are true e.g. if the effect was not foreseen but should have been.

role-oughts invoke ideals, and that (a) ideal-ought does not imply can. Indeed, when talking about teachers in the abstract it is natural to invoke ideals for teachers. (Recall Feldman’s “Teachers ought to explain things clearly”). And if ideal-ought does not imply can, we get the result that the incompetent teacher ought to do something that they are unable to:

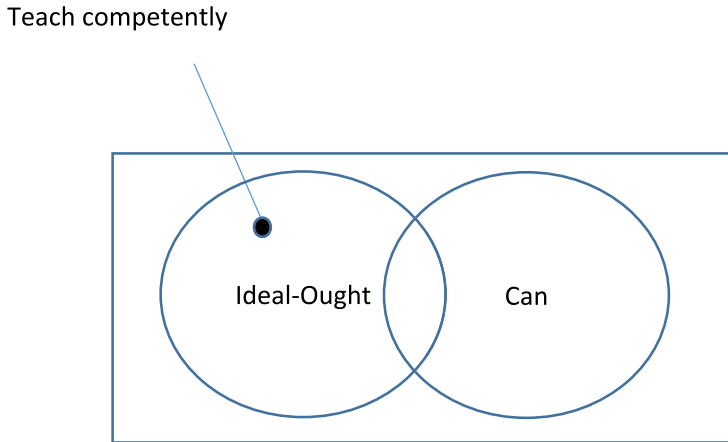


Figure 6: Ideally-ought to teach competently but cannot.

We can generate a different intuition if we add that the question is whether the teacher is blameworthy and fill in the case as one in which the teacher cannot explain things clearly through no fault of their own. Perhaps they are being asked to teach something that they have never been taught themselves. Then even if they ideally-ought to explain things clearly, we would not blame them for failing to explain things clearly, just as (c) predicts. In this hypothetical context inability excuses.

### 5.5. Addictive Behaviours

Recall that the argument from addictive behaviours says that in some cases agents ought to do things they are unable to do, thus refuting ought-implies-can.

The contextualist can explain this by understanding the context as invoking the ideal ought and using the hypothesis that (a) ideal-ought does not imply can. Moral or legal codes seem to express an ideal standard which does not take into account any limitations of the agent. So it is natural to understand moral and legal codes as expressing the ideal ought. Combined with (a) ideal ought does not imply can, it follows that the kleptomaniac ought not to steal (even if they can’t help themselves):

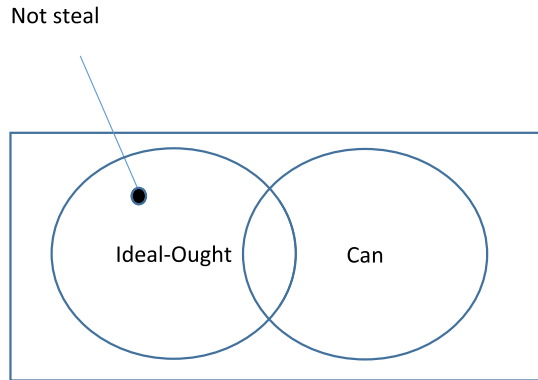


Figure 7: Ideal-ought not steal.

We can generate a different intuition if we use either the deliberative ought or hypothetical ought. Given that the agent cannot refrain from the addictive behaviour through no fault of their own, they do seem to have an excuse (as predicted by c) and should not deliberate about whether to refrain from the addictive behaviour (as predicted by b).

### 5.6 Fixed Feelings

Recall that the argument from fixed feelings says that agents ought to have feelings they are unable to have, thus refuting ought-implies-can.

The contextualist can explain this by understanding the context as invoking the ideal ought and the hypothesis that (a) ideal-ought does not imply can. The story involves the feelings that best fit the situation, where we do not take into account the limitations of the agent. The question is plausibly about the feelings of an ideal agent, i.e., the feelings one ideally-ought to have:




Figure 8: Ideal-ought to feel grateful.

Again, given that their feelings are fixed through no fault of their own, the agent should not deliberate about whether to have different feelings (as predicted by b) and does seem to have an excuse (as predicted by c).

This completes the main argument. I have argued that contextualism plus the three hypotheses above explain all the intuitions both for and against ought-implies-can. It remains to develop the view by responding to objections, tracing some dialectical consequences, and comparing alternatives.

## 6. Disagreement

A familiar objection is that contextualism wrongly predicts that there is no genuine disagreement. For example, according to contextualism the Professor and the student are using different parameters. Once the parameters are made explicit, each would accept what the other is saying. Disagreement disappears. And that looks like the wrong verdict, as there seems to be genuine disagreement. There are four well-known responses and I think there is some truth in all four, so I will briefly run through them. (I won't take a line on which response best applies to which cases).

A first response is that in many cases there isn't any disagreement. It is common for people to appear to disagree, only to find out that they have been using language slightly differently. Chalmers (2011) argues that this is typical even in philosophical debates. And if it happens in philosophical debates, it is all the more plausible that it happens in everyday conversations where speakers are far less reflective about the concepts they use.

A second response is that there is disagreement about *preferences over what to do* (Björnsson & Finlay 2010; Finlay 2014: ch. 8). The hidden goal of the conversation is to establish what to do. The Professor's preference is that students turn in their papers on time, and, failing that, that a just outcome is reached. The student's preference is that they get the highest mark with the minimum effort (we can assume). Thus there is no disagreement over which propositions are true, but disagreement over what to do.<sup>22</sup>

A third response is that disagreement involving contextualist terms should be understood as disagreement about *which parameters to use*. To motivate this view, note that contextualism says that the parameters are determined by the conversational context, and which parameters are operative might be in flux in a conversation. For example, consider the following conversation when all parties know that Michael Jordan is 1.98m tall and Sun Mingming is 2.45m tall:

---

22. For criticism, see McKenna (2014) and Bolinger (2022).

A: Michael Jordan is tall

B: No he isn't, Sun Mingming is tall

It is plausible that this conversation is best understood as an implicit negotiation about the standards for the word 'tall.'

Lewis (1979; 1996) suggests that conversations should conform to the principle of accommodation, which says that when a speaker makes an utterance involving a context-sensitive term, the parameters shift to make the sentence true. In Lewis's terms, the utterance changes the "conversational score." Thus B is attempting to shift the conversational score to one in which the relevant sense of "tall" is "tall-for-a-basketball player."

But what if the hearer does not want to change the conversational score? Then we get what Plunkett and Sundell (2013) call *metalinguistic negotiation*, in which, to continue the metaphor, each person tries to get their favoured values of parameters onto the scoresheet. Thus the apparent disagreement is really an instance of metalinguistic negotiation.<sup>23</sup> There is no guarantee that there will be a successful conclusion to the metalinguistic negotiation, resulting in an unresolved *disagreement*. This fits nicely with the sense of unresolved disagreement in the Professor and student case. The Professor is attempting to put a hypological context on the scoreboard, the student is attempting to put a deliberative context on the scoreboard, and neither is giving way to the other.

Someone might object that understanding a disagreement as a metalinguistic negotiation trivializes the disagreement, making it about parameters and concepts rather than about important matters like what we ought to do. But there is nothing trivial about this disagreement. As Bolinger writes:

While the dispute is *directly* over terms, it is indirectly over the fundamental structure of relevant moral concepts. (2022: 372, italics theirs)

More concretely, if you are being blamed and facing punishment for wrongdoing, then the question of which possibilities are live is anything but trivial.<sup>24</sup>

A fourth response is that there is a privileged 'ought', or an 'ought' with *authority*. The contextualist could add that speakers try to co-ordinate on the privileged 'ought' i.e. the correct values of parameters.<sup>25</sup> Even if speakers agree that the privileged ought is relevant, there can still be disagreement about which 'ought' that is. That is, there can still be disagreement about which are the correct parameters. In fact all that's needed for genuine disagreement is that the speakers have differing beliefs about which 'ought' is privileged (even if none of them are actually privileged).

23. See Horn (1989), Plunkett & Sundell (2013), and Bolinger (2022).

24. See also Stroud (2019).

25. See Worsnip (2019).



To be clear, it's not plausible that there is a privileged value of the aim parameter (express ideals, deliberate/guide, praise/blame) or the standard parameter (e.g. moral standard, practical standard). Those are up to the participants. But it is plausible that there is a privileged or correct set of live possible worlds *given* an aim and standard. For example, given a conversation about blame and the moral standard, it plausibly follows that some particular action (e.g. turning the paper in on time) is in the privileged or correct set of live possible worlds. So there is room for disagreement, and indeed for one party being right and another wrong, when it comes to the set of live possible worlds. This predicts that, once it is established that the conversation is about blame, the student cannot insist that there is no live possibility in which they hand in the paper.

## 7. Dialectical Consequences

I shall argue that the foregoing entails that ought-implies-can should not be used the way it often is in the literature.

Let's take a quick look at three influential examples. Howard-Snyder (1997) moves from ought-implies-can to the rejection of objective consequentialism.<sup>26</sup> The idea is that objective consequentialism requires us to perform the action with the best consequences, we often do not know what the consequences of our actions would be, so we cannot perform the action with best consequences. Assuming ought-implies-can, it follows that it's not the case that we ought to perform the action with best consequences.

Widerker (1991) argues that ought-implies-can, plus causal determinism, leads to the rejection of deontic claims such as "John's action was wrong."<sup>27</sup> The idea is that causal determinism entails that John could not have done anything else. As John could not have done anything else, ought-implies-can entails that it is not the case that he ought to have done something else. And if it is not the case that he ought to have done something else, then John's action is not wrong (nor is it right).

Alston (1988) argues from ought-implies-can to the rejection of deontic claims about beliefs such as "Bill's belief is unjustified."<sup>28</sup> The idea is that we do not have voluntary control over our beliefs, so John could not have believed something else. As Bill could not have believed something else, ought-implies-can entails that it is not the case that he ought to have believed something else. And if it is not the case that he ought to have believed something else then John's belief is not unjustified (nor is it justified).

---

26. See also Flanagan (1991).

27. See also Haji (1998; 2002).

28. See Helton (2020) for recent applications of ought-implies-can in epistemology.

These authors all identify a sense in which something cannot be done, then invoke ought-implies-can to infer that it is not the case that it ought to be done. But the context-sensitivity of 'ought' prevents such a quick inference.

In response to Howard-Snyder (1997), the fact that we don't know which action has the best consequences does not prevent objective consequentialism from being a theory of the ideal-ought. In response to Widerker (1991) the fact that John could not have done something else is compatible with that fact that he ideally-ought to have done something else. And in response to Alston (1988) the fact that Bill could not have believed anything else is compatible with that fact that he ideally-ought to have believed something else. The authors would need to argue that ought-implies-can is true in all relevant contexts. As utterances of the sentences the authors discuss could occur in many different contexts, it seems that no general argument could be given.

I do not say that no ought-implies-can principle is usable. But attention must be paid to the context to make sure there is no illicit shift of the values of parameters.

## 8. Alternative Weakenings of Ought-Implies-Can

I have defended a logical weakening of ought-implies-can: in some (not all) contexts, ought-implies-can. It might be useful to consider two other ways of weakening ought-implies-can.

Bassford (2022: 781) suggests making the consequent of ought-implies-can logically weaker:

[O]ught implies can or could have.

This does a nice job of dealing with cases like *Late Paper*. The student could have written the paper, so we retain the intuitive verdict that they ought to have done so.

But it cannot handle the arguments from role-oughts, addictive behaviours or fixed feelings. In cases where the agent cannot, nor could have, performed the action, it would follow that there is no obligation. But it seems that there is.

King (2019: 71) weakens ought-implies-can by making the antecedent logically stronger:

*Natural Ought Principle:*

If S naturally ought to A, then S can A.

*Non-Natural Ought Principle:*

It is not the case that, if S non-naturally ought to A, then S can A.

Very roughly, the natural oughts are those we have just by being human, while non-natural oughts are those we can only acquire voluntarily, e.g. by making a promise. King suggests that only the latter kind of obligations survive one's inability to perform the action.

But both principles are questionable. The student's obligation to hand in a paper is non-natural, but seems to survive his inability to do it. So that seems to be a case in which non-natural ought implies can. For a counterexample to natural-ought-implies-can, consider Aurelia's utterance "I ought to save him but can't get past the sharks." Suppose Aurelia has not willingly acquired any relevant obligations, e.g. she is not a lifeguard, she is just a passer-by. So we have a natural ought. Nevertheless, we can understand the utterance in a way that makes it true. Thus we seem to have a counter-example to natural-ought-implies-can.

## 9. Conclusion

There are powerful arguments both for and against the thesis that ought-implies-can. Rather than trying to explain away any intuitions it is better to make sense of the debate by explaining how opposing intuitions can be correct. One advantage of contextualism is that it does exactly that. Another advantage is that contextualism is independently motivated by work in linguistics. I have argued that contextualism can explain the intuition that ought-implies-can, while also predicting that in many contexts it is false that ought-implies-can. Specifically, I have defended three hypotheses:

- a) Ideal-ought does *not* imply can.
- b) Deliberative-ought *does* imply can.
- c) Hypological-ought *does* imply can, unless the agent is blameworthy for the inability

The main upshot is that philosophers cannot appeal to ought-implies-can in their arguments without paying attention to the context. However, it does follow that in deliberative contexts philosophers can use ought-implies-can, and that inability excuses (unless the agent is blameworthy for the inability).

## References

- Alston, William P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2, 257–299.
- Arpaly, N. (2003). *Unprincipled Virtue: An Inquiry into Moral Agency*. 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.
- Bassford, Andrew Dennis (2022). Ought Implies Can Or Could Have. *Review of Metaphysics* 75(4), 779–807.
- Berg, Amy (2018). Ideal Theory and ‘Ought Implies Can’. *Pacific Philosophical Quarterly*, 99(4), 869–890.
- Björnsson, Gunnar, and Stephen Finlay (2010). Metaethical Contextualism Defended. *Ethics*, 121, 7–36.
- Bolinger, Rene J. (2022). Metalinguistic negotiations in moral disagreement. *Inquiry*, 65(3), 352–380.
- Bradley, Darren (2021). Ought-contextualism and reasoning. *Synthese*, 199(1–2), 2977–2999.
- Brogaard, Berit (2008). Moral Contextualism and Moral Relativism. *Philosophical Quarterly*, 58(232), 385–409.
- Buckwalter, Wesley (2020). Theoretical Motivation of ‘Ought Implies Can’. *Philosophia*, 48, 83–94. <https://doi.org/10.1007/s11406-019-00083-7>
- Cariani, Fabrizio, Magdalena Kaufmann, and Stefan Kaufmann (2013). Deliberative Modality under Epistemic Uncertainty. *Linguistics and Philosophy*, 36(3), 225–259.
- Carr, Jeniffer (2015). Subjective Ought. *Ergo: An Open Access Journal of Philosophy*, 2(27), 678–710.
- Chalmers, David J. (2011). Verbal Disputes. *Philosophical Review*, 120(4), 515–566.
- Charlow, Nate (2013). What we know and what to do. *Synthese*, 190(12), 2291–2323.
- Chrisman, Matthew (2012). On the Meaning of ‘Ought’. *Oxford Studies in Metaethics*, 7, 304–332.
- Chrisman, M. (2015). *The Meaning of ‘Ought’: Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press.
- Chuard, Philippe and Nicholas Southwood (2009). Epistemic Norms Without Voluntary Control. *Noûs*, 43(4), 599–632.
- Copp, David (1995). *Morality, Normativity, and Society*. Oxford University Press.
- Copp, David (1997). Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility. *Noûs*, 31(4), 441–456.
- Copp, David (2008). ‘Ought’ implies ‘can’ and the derivation of the Principle of Alternate Possibilities. *Analysis* 68(1), 67–75.
- Dowell, Janice (2012). Contextualist Solutions to Three Puzzles about Practical Conditional. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics Vol. 7* (271–303).
- Dowell, Janice (2013). Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity. *Inquiry*, 56(2–3), 149–178.
- Dreier, Jamie (1990). Internalism and Speaker Relativism. *Ethics*, 101, 6–26.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Fischer, John M. (2003). ‘Ought-implies-can’, causal determinism and moral responsibility.” *Analysis* 63(3), 244–250.
- Flanagan, Owen (1991). *Varieties of Moral Personality*. Harvard University Press.
- Foot, P. (1972). Morality as a System of Hypothetical Imperatives. *The Philosophical Review*, 81(3), 305–316.
- Geach, Peter (1956). Good and evil. *Analysis*, 17(2), 33–42.
- Graham, Peter A. (2011). ‘Ought’ and Ability. *Philosophical Review*, 120(3), 337–382.
- Haji, Ishtiyaque (1998). *Moral Appraisability*. Oxford University Press.

- Haji, Ishtiyaque (2002) *Deontic Morality and Control*. Cambridge University Press.
- Hare, Richard M. (1963). *Freedom and Reason*. Oxford University Press.
- Harman, Gilbert (1975). Moral Relativism Defended. *The Philosophical Review*, 84(1), 3–22.
- Hedden, Brian (2012). Options and the subjective ought. *Philosophical Studies*, 158(2), 343–360.
- Helton, G. (2020). If You Can't Change What You Believe, You Don't Believe It. *Noûs*, 54(3), 501–526.
- Henne, Paul, Jennifer Semler, Vladimir Chituc, Felipe De Brigard, and Walter Sinnott-Armstrong (2019). Against Some Recent Arguments for 'Ought' Implies 'Can': Reasons, Deliberation, Trying, and Furniture. *Philosophia*, 47(1), 131–139
- Heuer, Ulrike (2010). Reasons and impossibility. *Philosophical Studies*, 147(2), 235–246.
- Hicks, Amelia (2022). Dispensing with the Subjective Moral 'Ought'. *Oxford Studies in Normative Ethics*, 11. Oxford University Press.
- Horn, Laurence R. (1989). *A Natural History of Negation*. University of Chicago Press.
- Howard-Snyder, Frances (1997). The rejection of objective consequentialism. *Utilitas*, 9(2), 241–248.
- Howard-Snyder, Frances (2006). "Cannot" Implies "Not Ought." *Philosophical Studies*, 130(2), 233–246.
- Khoo, Joshua and Justin Knobe (2018). Moral Disagreement and Moral Semantics. *Noûs*, 52(1), 109–143.
- King, Alex (2019). *What We Ought and What We Can*. Routledge.
- Kolodny, Niko and John MacFarlane (2010). Ifs and Oughts. *Journal of Philosophy*, 107(3), 115–143.
- Kratzer, Angelika (1981). The Notional Category of Modality. In Hans-Jürgen Eikmeyer and Hannes Rieser (Eds.), *Words, Worlds, and Contexts: New Approaches in World Semantics* (38–74). de Gruyter.
- Kratzer, Angelika (1991). Modality. In Arnim von Stechow and Dieter Wunderlich (Eds.), *Semantics: An International Handbook of Contemporary Research* (639–650). de Gruyter.
- Kratzer, Angelika (2012). *Modals and Conditionals*. Oxford University Press.
- Lewis, David (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, David (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- McHugh, C. (2012). The Truth Norm of Belief. *Pacific Philosophical Quarterly*, 93(1), 8–30.
- McKenna, Robin (2014). Disagreeing About 'Ought'. *Ethics*, 124(3), 589–597.
- Miller, D. J. (2017). Reasonable foreseeability and blameless ignorance. *Philosophical Studies*, 174, 1561–1581.
- Moore, G. E. (1922). *Philosophical Studies*. Keegan Paul, Trench, Trubner, & Co.
- Norcross, Alastair (2020). *Morality by Degrees: Reasons without Demands*. Oxford University Press.
- Plunkett, David and Timothy Sundell (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, 13(23): 1–37.
- Rosen, Gideon (2004). Skepticism about Moral Responsibility. *Philosophical Perspectives*, 18(1), 295–313.
- Ryan, Sharon (2003). Doxastic Compatibilism and the Ethics of Belief. *Philosophical Studies*, 114(1), 47–79.
- Schoenfield, Miriam (2018). An Accuracy Based Approach to Higher Order Evidence. *Philosophy and Phenomenological Research*, 96(3), 690–715.
- Schroeder, Mark (2011). *Ought, Agents, and Actions*. *Philosophical Review*, 120(1), 1–41.

- Schwarz, Wolfgang (2020). Ability and Possibility. *Philosophers' Imprint*, 20(6).
- Silk, Alex (2016). *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford University Press.
- Silk, Alex (2017). Normative Language in Context. *Oxford Studies in Metaethics*, 12, 206–239.
- Sinnott-Armstrong, Walter (1984). "Ought" Conversationally Implies "Can." *The Philosophical Review*, 93(2), 249–261.
- Steinberger, F. (2019). Three ways in which logic might be normative. *The Journal of Philosophy*, 116(1), 5–31.
- Stocker, Michael (1971). 'Ought' and 'Can'. *Australasian Journal of Philosophy*, 49(3), 303–316.
- Streumer, Bart (2003). Does 'Ought' Conversationally Implicate 'Can'? *European Journal of Philosophy*, 11(2), 219–228.
- Streumer, Bart (2007). Reasons and Impossibility. *Philosophical Studies*, 136(3), 351–384.
- Stroud, Sarah (2019). Conceptual disagreement. *American Philosophical Quarterly*, 56(1), 15–28.
- Suikkanen, Jussi (2020). Consequentializing Moral Dilemmas. *Journal of Moral Philosophy*, 17(3), 261–289.
- Talbot, Brian (2021). Why Impossible Options Are Better: Consequentializing Dilemmas. *Utilitas*, 33(2), 221–236.
- Vetter, Barbara and Emanuel Viebahn (2016). How many meanings for 'may'? The case for modal polysemy. *Philosophers' Imprint*, 16(10).
- Vranas, Peter B. (2007). I Ought, Therefore I Can. *Philosophical Studies*, 136(2), 167–216.
- Vranas, Peter B. (2018). "Ought" Implies "Can" but Does Not Imply "Must": An Asymmetry between Becoming Infeasible and Becoming Overridden. *The Philosophical Review*, 127(4), 487–514.
- Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford University Press.
- Widerker, David (1991). Frankfurt on 'Ought implies Can' and alternative possibilities. *Analysis*, 51(4), 222–224.
- Williams, Bernard (1979). Internal and External Reasons. In H. Ross (Ed.), *Rational Action* (101–113). Cambridge University Press.
- Worsnip, A. (2019). 'Ought'-contextualism beyond the parochial. *Philosophical Studies*, 176(11), 3099–3119.
- Zimmerman, Michael (1996). *The Concept of Moral Obligation*. Cambridge University Press.
- Zimmerman, Michael (1997). Moral Responsibility and Ignorance. *Ethics*, 107(3), 410–426.