

# MOOREAN PROMISES

Bob Beddor

Forthcoming in *Ethics*

## Abstract

“I promise to mow your lawn, but I don’t know whether I will.” Call promises of this form, “Moorean”, based on their resemblance to Moore’s paradox. Moorean promises sound absurd. But why? In the literature on assertion, many have used Moore’s paradox to motivate a knowledge norm of assertion. I put forward an analogous norm on promising, according to which one should only make a promise if one knows one will fulfill it. A knowledge norm explains why Moorean promises are absurd; it accounts for a variety of linguistic data; and it sheds light on how promises generate obligations.

## 1 A Strange Sort of Promise

You mention to your neighbor that you’re going on vacation. They say:

- (1) # I promise to walk your dog while you’re gone, but I don’t know whether I will.

You would regard their promise as absurd. More generally, there seems to be something incoherent about making a promise while simultaneously confessing ignorance as to whether you will carry through. A smattering of examples:

- (2) # I’ll water your plants, but I might not do so.
- (3) # I don’t know if I’ll make it to your recital, but I promise to attend.
- (4) # I promise to meet you at the theater at 7pm, though I wonder if I’ll be able to make it to the theater by then.

All of these “Moorean promises” sound terrible. Why?

The answer does not hinge on any quirks of the English verb “promise”, for two reasons. First, it’s a familiar point that someone can make a promise without intoning “I promise.” In the right context, I can promise to water your plants by simply claiming that I *will* water them. But it still sounds absurd to conjoin this promise with an admission that I might not comply, as revealed by (2). Second, these Moorean data are cross-linguistically robust. Here are some examples from other languages, representing four different language families (Austroasiatic, Indo-European, Sino-Tibetan, and Japonic):

- (5) # Tôi hứa sẽ dẫn chó của bạn đi dạo nhưng tôi không biết mình có làm (hay) không. **(Vietnamese)**  
 “I promise to walk your dog, but I don’t know whether I will.”
- (6) # Ek below om jou plante water te gee, maar ek sal dit dalk nie doen nie. **(Afrikaans)**  
 “I promise to water your plants, but I might not do it.”
- (7) # Obecavam da cu da uradim to, ali ne znam da li cu. **(Serbian)**  
 “I promise to do it, but I don’t know whether I will.”
- (8) # Wǒ dāyìng bāng nǐ liú gǒu, dàn wǒ bù quèdìng huì háishì bù huì. **(Mandarin)**  
 “I promise to help you walk the dog, but I am not certain whether I will.”
- (9) # (Ashita) inu-o sampo su-ru to yakusoku su-ru kedo, (hontoo-ni) su-ru kadooka wakara-nai. **(Japanese)**  
 “I promise that I will walk the dog tomorrow, but I don’t know whether I will or not.”

Speakers of these languages find these Moorean promises just as degraded as their English counterparts.<sup>1</sup> So the Moorean data are not just artifacts of English. Rather, they reveal something general about the conditions under which it is appropriate to make a promise.

If we turn to the extensive literature on promising, we will be hard-pressed to find an explanation for why Moorean promises are absurd.<sup>2</sup> Many philosophers have held that promises convey an intention to perform the promised action.<sup>3</sup> However, we’ll see that this view is too weak to explain the absurdity of Moorean promises. I might rationally intend to make it to your recital, even though I don’t know that I will be able to make it, and even though I know that I do not know this. But in this situation, making a Moorean promise such as (3) sounds no less absurd.

Another major theme in the literature on promising is that promises voluntarily create obligations and bestow rights.<sup>4</sup> When your neighbor promises to walk your dog, they incur an obligation to walk your dog. They also give you the right to release them from this obligation. These observations—while surely correct—do nothing to explain why Moorean promises are absurd. Why can’t your neighbor incur an obligation to walk your dog while also acknowledging that they do not know whether they will fulfill this obligation?

<sup>1</sup>Thanks to Anne Nguyen, Andries Coetzee, Jelena Krivokapic, and Mitcho Erlewine for these examples and their judgments. Thanks also to Savi Namboodiripad for further discussion of the cross-linguistic data.

<sup>2</sup>Some philosophers explicitly argue that it can be appropriate to make a promise in full awareness that one may not fulfill one’s promise (Marušić 2013, 2015; Liberman 2019). Such views have particular trouble explaining the absurdity of Moorean promises. I discuss these arguments in detail in §8.

<sup>3</sup>This idea traces back to Hume, who remarked that “when a man says he promises anything, he in effect expresses a resolution of performing it” (Hume 1739/1978: 522). The connection between promising and intending also has many contemporary advocates; see Austin 1962; Searle 1969; Bach and Harnish 1979; Scanlon 1990; Marušić 2013, 2015, among others. Arguably, some version of this connection is also reflected in contract law. For example, the Restatement (Second) of Contracts characterizes a promise in terms of “a manifestation of an intention to act or refrain from acting in a specific way” §2 (Am. Law Inst. 1981).

<sup>4</sup>This tradition also has its roots in Hume, who claimed that “every promise creates a new obligation on the person who promises” (Hume 1739/1978: 524). See also Anscombe 1978; Raz 1972, 1977; Prichard 2002; Owens 2006a, 2014, among many others.

While views that emphasize the normative powers of promises do not explain the Moorean data, the order of explanation may go the other way around: it may turn out that the correct explanation of the Moorean data also sheds light on the normative properties of promises. Or so I shall argue.

## 2 The Path Ahead

Since mainstream views of promising fail to explain the Moorean data, this paper develops a more promising approach. My approach starts by taking seriously the parallels between Moorean promises and their more well-known cousins, Moorean assertions, e.g.:

- (10) # You have a dog, but I don't know whether you have a dog.
- (11) # Your plants need to be watered, but they might not need to be watered.
- (12) # I don't know if your recital is tonight, but your recital is tonight.
- (13) # The movie starts at 7pm, though I wonder whether the movie starts at 7pm.

One leading solution to Moore's paradox holds that knowledge is the norm of assertion. According to this approach, one should only assert  $p$  if one knows  $p$ . Moorean assertions are defective because they are guaranteed to violate this norm. This paper develops a parallel treatment of Moorean promises, using an epistemic norm on promising. According to this norm, one should only promise to  $\phi$  if one knows that one will  $\phi$ .<sup>5</sup>

In addition to explaining the Moorean data, this epistemic norm on promising carries a number of downstream benefits. It sheds light on the place of promises in the speech act landscape, providing fresh support for the view that promises are a species of assertion, and revealing some surprising differences between promises and imperatives. It also provides new insight into how promises generate obligations. Specifically, this epistemic norm offers an elegant solution to the circularity problem for expectation-based views of promissory obligation.

## 3 Moorean Assertions

Why is it absurd to assert  $p$  while acknowledging that you don't know  $p$ ? By now, a number of answers have been proposed. I will focus on one particularly compelling explanation, which traces back to Moore's own remarks about his paradox. According to Moore, sentences like (10) carry an important lesson about the nature of assertion: "by asserting  $p$  positively, you imply, though you don't assert, that you know that  $p$ " (Moore 1962: 277).

Williamson 2000 famously cashes out Moore's insight in terms of a normative constraint on assertion:

---

<sup>5</sup>This is not to deny that it's possible for someone to make a promise without knowing whether they will fulfill it. A faulty promise—that is, a promise that runs afoul of the norms on promising—is still a promise, just as one can still assert that which one does not know. Throughout this paper, my primary focus will be on the *norms* of promise-making.

KNOWLEDGE NORM OF ASSERTION (KA) One ought only assert  $p$  if one knows  $p$ .<sup>6</sup>

As Williamson 2000 observes, KA explains the absurdity of Moorean assertions. A distinctive feature of Moorean assertions is that they are *unknowable*. Take (10) (“You have a dog, but I don’t know whether you have a dog”). If the speaker knows the first conjunct, then the second conjunct is false, and so cannot be known. So anyone who utters (10) is guaranteed to violate the norm of assertion.

This explanation generalizes to other Moorean absurdities. Take (11) (“Your plants need to be watered, but they might not need to be watered”). Suppose we make the plausible assumption that knowing  $p$  is incompatible with knowing that  $p$  might be false ( $Kp \Rightarrow \neg K\Diamond\neg p$ ). Then (11) is unknowable, and so guaranteed to violate KA. This explanation also extends to (13) (“The movie starts at 7pm, though I wonder whether the movie starts at 7pm”), given the plausible premise that it is impossible—or at least rationally defective—to wonder whether  $p$  while at the same time knowing  $p$ .<sup>7</sup>

Some might wonder whether KA explains the full range of Mooreanisms. Consider:

(14) # You teach a seminar on Monday, but I’m not certain you teach on Monday.

(14) also seems infelicitous (Unger 1971; Stanley 2008). Does KA explain why?

There are at least two possible responses. One is to follow Williamson’s suggestion that, in ordinary conversation, we are reluctant to let the “contextually set standards for knowledge and certainty diverge” (2000: 204). According to this response, we tend to infer a knowledge denial from a certainty denial, so (14) is infelicitous for the very same reason as (10). The second option is strengthen our norm of assertion. Perhaps one ought only assert  $p$  if one knows  $p$  with certainty (Stanley 2008; Beddor 2020a,b; Goodman and Holguín forthcoming). For our purposes, we can remain noncommittal on which of these two responses is preferable.<sup>8</sup>

What sort of norm is KA? This is debated in the literature. But most agree with at least the following: KA is an epistemic norm regulating the speech act of assertion. Specifically, it tells us what sort of epistemic condition is required for someone to have the authority to make an assertion (Williamson 2000: 257). As such, it is both an epistemic norm and a conversational norm. But unlike some other conversational norms, it is *strict*. Speakers cannot simply opt of following the norm (otherwise Moorean assertions would not sound so bad after all). Williamson goes a step further, arguing that it is a constitutive norm of assertion. Just as the rules of chess partially determine what counts as castling, so too

---

<sup>6</sup>Similar ideas are defended by Unger 1975; Slote 1979; DeRose 2002; Schaffer 2008; Benton 2011; Turri 2011; Kelp 2018, among others.

<sup>7</sup>For defense of this premise, see Friedman 2013; Sapir and van Elswyk 2021.

<sup>8</sup>If certainty is seldom attainable, as many philosophers have contended, this might seem to impose an impossibly demanding constraint on assertion. However, in everyday discourse we frequently ascribe certainty to ourselves and others—e.g., “Scientists are absolutely certain that this warming trend is due to human activity.” (<https://www.brookings.edu/wp-content/uploads/2019/09/20190920-global-response-to-the-climate-crisis.pdf>) For development of a contextualist treatment of certainty aimed at accommodating this observation, see Beddor 2020a.

KA partially determines what it takes for a speech act to count as an assertion. For our purposes, we can remain neutral on whether KA is constitutive in this sense.

## 4 Knowledge and Promising

The basic idea behind my proposal is simple. We should extend our norm on assertion to encompass promising:

KNOWLEDGE NORM OF PROMISING (KP): One ought not promise to  $\phi$  unless one knows one will  $\phi$ .

A couple points of clarification are in order. First, what sort of norm is KP? The short answer: whatever sort of norm KA is! Just as KA is an epistemic norm regulating the speech act of assertion, KP is an epistemic norm regulating the speech act of promising. As such, it is both an epistemic norm and a conversational norm. But, much like KA, it is stricter than many conversational norms; one cannot simply opt out of it. Fans of constitutive norms could go a step further, and hold that KP partially determines what it is for a speech act to be a promise. Here too, we can afford to be agnostic on whether KP is constitutive.

Second, KP does not require that a promisor knows that they will keep their promise *before* promising. Such a requirement would run afoul of spontaneous promises: I might make a spur of the moment promise to attend your recital, even though before opening my mouth I had no plan to attend. Rather, KP should be understood as forbidding one from making a promise unless one knows *at the time the promise is made* that one will keep it. As soon as I blurt out my promise to attend your recital, I would be *prima facie* criticizable if I did not take myself to know I would attend.

The situation with assertion is no different. It can be perfectly fine to assert:

(15) I am speaking now.

But the speaker did not know the content of (15) before the time of utterance, since before the time of utterance they were not speaking. Is this a counterexample to KA? Hardly. It only shows that we should construe KA as requiring that a speaker knows  $p$  at the time of asserting  $p$ , not some time beforehand.

With these clarifications in place, let us turn to the main question: why accept KP? An initial argument comes from the tight link between one's position to promise and one's position to assert:

PROMISE-ASSERTION BRIDGE Anyone who is in a position to promise to  $\phi$  is in a position to assert that they will  $\phi$ .

This principle seems very plausible. Compare:

(16) I'll come to your recital.

(17) I promise that I'll come to your recital.

Intuitively, (17) makes at least as strong a claim as (16).<sup>9</sup>

Combining Promise-Assertion Bridge with KA leads to KP. By Promise-Assertion Bridge, you are only in a position to promise to attend the recital if you are in a position to assert that you will attend. By KA, you are only in a position to assert that you will attend if you are in a position to know you will do so. So you are in a position to promise to attend the recital only if you are in a position to know that you will attend.

This was just a warm-up. The main argument for KP is that it explains the absurdity of Moorean promises. Now, some might worry that there is a subtle difference between Moorean assertions and Moorean promises—a difference that make the latter more difficult to explain. When it came to Moorean assertions, we relied on the fact that both conjuncts are asserted in order to explain why Moorean assertions violate KA. But consider a Moorean promise such as (1) (“I promise to walk your dog, but I don’t know whether I will”). Only the first conjunct is a promise. The second conjunct (“I don’t know whether I will”) is an assertion. So we cannot use KP to show that Moorean promises are guaranteed to violate the norm of promising.

But if we take on board KA, we can address this worry. If your neighbor asserts (1), either they know that they will walk your dog or they don’t know this. Suppose they know it. Then while they have not violated KP, the second conjunct of (1) is false, hence violates KA. Suppose then they don’t know they will walk your dog. Then the first conjunct violates KP. The upshot: Moorean promises are guaranteed to violate the norms governing at least one of the speech acts involved—either they violate the norm of assertion (KA) or they violate the norm of promising (KP).

This explanation extends to other Moorean promises. Moorean promises involving epistemic modals such as (2) (“I promise to water your plants, but I might not do so”) are also guaranteed to violate either KA or KP, given the assumption that knowing  $\neg p$  precludes knowing  $p$ . Similarly, the absurdity of (18) (“I promise to meet you at the theater at 7pm, though I wonder if I’ll be able to make it to the theater by then”) follows from KA and KP, given the assumption that knowing  $p$  precludes wondering whether  $p$ .

As with KA, some might wonder whether KP explains the full range of Moorean absurdities. There seems to be something absurd about promising to do something while professing to be less than certain about whether you will carry through:

(18) # I promise to teach your Monday seminar, but I’m not certain I will.

Here too, we can explain this infelicity by either embracing the idea that knowledge entails certainty, or by replacing KP with a stronger norm, according to which one ought only promise to  $\phi$  if one knows with certainty that one will  $\phi$ . In short, the same maneuvers available to KA are equally available to KP.

---

<sup>9</sup>Thanks to Harvey Lederman for this point.

## 5 Further Support for KP

The case for a knowledge norm of assertion does not rely exclusively on Moorean data. Here I consider some auxiliary arguments that have been put forward on its behalf. I argue that each can be adapted to bolster a knowledge norm of promising. While these auxiliary arguments may not, when examined one-by-one, prove decisive, taken together, they mount a compelling abductive case for KP.

### 5.1 I Can't Promise That

When invited to answer some question, we sometimes “opt out” by confessing we do not *know* the answer (Dorst 2014; Benton 2023). For example:

- (19) a. *A*: How do you get to Canal Street from here?  
b. *B*: Sorry, I can't say—I don't know where Canal Street is.

If assertion was not governed by a knowledge norm, it is hard to see why *B*'s lack of knowledge gets them off the hook for providing an answer.

Promises exhibit a similar pattern. When invited to make a promise, we sometimes opt out by saying, “I can't promise that.” Moreover, the reason for opting out is often that we do not *know* we will make good on the promise. For an illustration “from the wild”, consider the following extended excerpt from *I promise*, a *That 70s Show* fan fic origin story in which Red and Kitty have a lachrymose *tête-à-tête* before Red goes off to war:

“Will you promise you'll come back home to me?” she asked, her voice trembling with emotion.

“Kitty...”

“I know, I know. You can't promise that. No one can make that kind of promise right now.” Kitty shook her head [...] “Will you promise to write me every day?”

Red gently used his left thumb to stroke her hand he was holding. “Some days are busier than others. Some days I'll be lucky to get a night's sleep.”

“Promise me you'll be careful then.”

“War's not exactly a careful place, Kitty,” he sadly reminded her [...]

“Red, honey, I'm... I'm trying to hold on to some kind of hope here,” Kitty was trying to keep her emotions in check but looked like she was on the verge of bursting into tears. “Isn't there anything you can promise me?” [...]

“There is so much I want to promise you, sweetheart. I wanna promise that when I get back home, we'll get married and start our family. [...] *I wanna promise you the world, Kitty*. But I can't make any of those promises, knowing that there's a chance I might not be able to keep them.”<sup>10</sup>

<sup>10</sup>Source: <https://www.fanfiction.net/s/13343717/1/Magic-Moments-Filled-with-Love>

Some might object that there are other possible explanations for why Red refuses to make the promises that Kitty solicits. Perhaps Red cannot promise to return, to write daily, or to be careful not because he doesn't *know* he will, but rather because these things are either not likely, given his evidence (writing daily, being careful), or not under his control (returning from war).<sup>11</sup> However, these alternative explanations do not do justice to Red's stated reason for why he cannot make any of these promises. Consider the last line of the dialogue. Red's reason is not that he is unlikely to do these things, or that they are out of his control. Rather, it is something weaker: *he knows there's a chance that he might not be able to keep these promises*. Why does knowing there's a chance he might not be able to keep these promises provide a sufficient justification for refraining from making them? KP provides an answer. Knowing one might not be able to  $\phi$  entails not knowing one will  $\phi$ . Consequently, Red cannot make these promises while abiding by KP.

For another example, consider the following excerpt from the negotiations between Lê Đức Thọ and Henry Kissinger. In their conversation on February 13, 1973, Lê Đức Thọ repeatedly presses Kissinger to commit to a shorter timeline for landmine removal. Kissinger balks, insisting that landmine removal is a complicated and time-consuming process, and that he doesn't know whether it can be done more quickly:

*Lê Đức Thọ*: You should promise that it will be done with greater speed.

*Kissinger*: I can't promise what I don't know.<sup>12</sup>

As these examples illustrate, people often treat their lack of knowledge of whether they will  $\phi$  to be sufficient grounds for refusing to promise to  $\phi$ . This is exactly what we should expect if knowledge is the norm of promising. By contrast, weaker constraints on promising struggle to explain this phenomenon. If permissible promising merely requires (say) a high degree of confidence that one will fulfill one's promise, or that the promised outcome is under one's control, why should one's lack of *knowledge* get one off the hook for making a promise?

## 5.2 Conversational Presuppositions

Another argument for KA comes from the fact that asserting  $p$  typically licenses the presupposition that the speaker knows  $p$  (Unger 1975: 263-265; Williamson 2000: 252-253; Turri 2010; Benton 2023):

- (20) a. *A*: The bus is going to be late.  
 b. *B*: How do you know that?  
 c. *A*: # Sheesh, I didn't say I *knew* the bus was gonna be late. I just said it was gonna be late!

*B*'s question presupposes that *A* knows the bus will be late. This presupposition seems warranted, as shown by the strangeness of *A*'s reply. But why is it warranted? Here is

<sup>11</sup>Thanks to a referee for raising this point.

<sup>12</sup>Source: <https://history.state.gov/historicaldocuments/frus1969-76v42/d56>



an attractive explanation. We have a default entitlement to expect that our interlocutors abide by conversational norms:

EXPECTED COMPLIANCE When someone performs a particular speech act, we are defeasibly entitled to expect that they satisfy the norms governing the speech act.<sup>13</sup>

If knowledge is the norm of assertion, it follows that *B* is entitled to assume that *A* knew their assertion is true.

Here too, this argument carries over to promises. Consider a less moving variant of Red and Kitty's teary dialogue:

- (21) a. *Red*: I promise I'll be home by Christmas.  
b. *Kitty*: How do you know you'll be back by then?  
c. *Red*: # Sheesh, I didn't say I *knew* I would be back by then. I just promised to be back by then!

Kitty's question presupposes that Red knows that he will be home by Christmas. This presupposition seems warranted, as revealed by the inappropriateness of Red's reply. But why is this presupposition warranted? The story goes much the same. When someone makes a promise, we are defeasibly entitled to assume that they satisfy the norms governing promising. If knowledge is the norm of promising, it follows that Kitty is entitled to assume Red knew he would comply with his promise.

Some might question how far this argument extends. According to [Marušić 2015](#), many promises do *not* license knowledge presuppositions:

[W]hen I tell my spouse, "I will run twice a week from now on," it would be inappropriate or odd for her to ask, "How do you know?" This suggests that when we promise to do something, or express our resolution to do it, we don't necessarily represent ourselves as speaking from knowledge. ([Marušić 2015](#): 152)

However, I think this conclusion is too quick. In many contexts, it will be odd to respond to a promise with, "How do you know?" simply because the answer is obvious. In [Marušić's](#) example, the most likely reply would be something along the lines of, "Well, because I'm going to try to run twice a week, and I know I will do it if I try"—an answer so obvious that it is strange to be asked to articulate it.<sup>14</sup> Here the situation with respect to assertion

---

<sup>13</sup>A version of this principle lies at the heart of Gricean pragmatics ([Grice 1975](#)). According to Grice, we are defeasibly entitled to expect that our interlocutors are abiding by the conversational maxims (or at least the cooperative principle), and this expectation is what leads us to compute implicatures. For example, a characteristic line of Gricean reasoning might run: "My interlocutor said *p*, but this is only relevant to the conversation if *q* is true. Since I assume my interlocutor is abiding by the maxim of relevance, they must believe *q*."

<sup>14</sup>Of course, we can add details that will cast doubt on the adequacy of this answer. Maybe the promisor has sprained their ankle, calling into question whether they really know that they will run twice a week if they try. But once we add these details, it becomes more appropriate (even if potentially undiplomatic) to ask how the promisor knows they will carry through.

is much the same. If in the afternoon I assert I will want to have dinner in a few hours, it would be rather odd to ask, “How do you know?” Here the most likely reply would be something along the lines of, “Well, in a few hours it will be dinnertime, and I tend to get hungry if I don’t have dinner”—an answer so obvious that would be odd to have someone spell it out.

### 5.3 Lottery Promises

You have a single ticket in a lottery. The odds of winning are bleak—just one in ten thousand. Still, [Williamson 2000](#) observes, there seems to be something inappropriate about flat-out asserting that your ticket won’t win. KA explains why. You do not know that your ticket will lose. Indeed, you are in a position to know that you do not know this.

Tweak this example and we get an argument for KP. You are participating in a lethal lottery, where the prize is your demise. Each contestant drops their name in a hat. One name will be drawn, and this “winner” faces death (think the Shirley Jackson story, or *The Hunger Games*). Before dropping your name in the hat, you reassure your concerned family:

(22) # I promise I won’t win.

Your promise seems inappropriate. It would be natural for your loved ones to object, “You’re in no position to promise that!” KP explains why they are justified in making this protest. You are in no position to make your promise because you don’t know that your name will not be drawn. Indeed, you are in a position to know that you don’t know this.

This is not to say KP is the only possible explanation for why your lottery promise is inappropriate. Some might argue that losing the lottery is not the sort of action that is under your control, or that you can rationally intend. Perhaps it is these features—rather than your lack of knowledge—that explains why (22) is defective. I will defer a fuller discussion of this issue to §9, where I consider in detail some alternative explanations of our data. For now, I want to emphasize that KP delivers a unified and parsimonious explanation of all the data discussed so far. Specifically, it explains why (i) Moorean promises are infelicitous, (ii) one can abstain from making a promise by citing one’s lack of knowledge of whether one will carry through, (iii) promising to  $\phi$  typically licenses the presupposition that the promisor knows they will  $\phi$ , (iv) lottery promises are defective. As we will see in §9, the most natural alternative explanations either fail to account for the full range of data or they collapse into the knowledge norm.

## 6 The Function of Promising

So far, I have argued for KP primarily on the basis of linguistic data. But we can also argue for KP on non-linguistic grounds by considering the function of promising.

What’s the point of making a promise? According to one prominent view, to make a promise is to invite the promisee to rely on the promised action.<sup>15</sup> While it is con-

---

<sup>15</sup>For versions of this “Reliance View”, see [MacCormick 1972](#); [Thomson 1990](#); [Heape 2022](#).

troverial whether every promise invites reliance,<sup>16</sup> it seems hard to dispute that this a characteristic function of promising. As Heape notes, it would be normal to protest a broken promise by complaining, “But I was relying on you!” (Heape 2022: 151). The idea that promises invite reliance also sheds light on why it is wrong to break promises: it is *pro tanto* wrong to break a promise because in doing so you might be letting down someone who relied on your word.<sup>17</sup> The idea that promises invite reliance is also reflected in contract law. For example, the doctrine of promissory estoppel allows that when *A*’s promise foreseeably induces *B*’s reliance in a way that results in *B*’s detriment, then *B* can recover damages from *A* even when that promise would otherwise be unenforceable.<sup>18</sup>

This aspect of promising also has a normative dimension. Typically, promises do not just cause reliance, they *warrant* it. If you promise to mow my lawn, then, *ceteris paribus*, it would be epistemically appropriate for me to rely on the promise that you will mow my lawn. That is:

PROMISES WARRANT RELIANCE A characteristic function of promising to  $\phi$  is to warrant the promisee in relying on the premise that the promisor will  $\phi$ .

What sort of epistemic position does one need to bear to a proposition in order to be warranted in relying on it? According to one prominent view in epistemology, the answer is *knowledge*:

KNOWLEDGE NORM OF RELIANCE An agent *A* is warranted in relying on *p* in practical and theoretical reasoning if and only if *A* knows *p*.<sup>19</sup>

A norm along these lines gains support from our everyday patterns of criticism. As Hawthorne and Stanley observe: “If a parent allows a child to play near a dog and does not know whether the dog would bite the child, and if a doctor uses a needle that he did not know to be safe, then they are *prima facie* negligent” (2008: 572). Further support comes from reasoning about lotteries. If you have a ticket with a 1/10,000 chance of winning, it seems irrational to rely on the premise that it will lose—for example, by throwing it away. The Knowledge Norm of Reliance explains why: even though you have a justified (and perhaps true) belief that your ticket will lose, this belief does not amount to knowledge (Hawthorne 2004; Hawthorne and Stanley 2008).<sup>20</sup>

<sup>16</sup>See e.g., Friedrich and Southwood 2011 for doubts on this score.

<sup>17</sup>This is not to say that it wouldn’t be *pro tanto* wrong to break a promise that no one relies upon. At the very least, a promise that you *think* might be relied upon creates an obligation. I revisit these issues in §10.

<sup>18</sup>See Restatement (Second) of Contracts, §90.

<sup>19</sup>See e.g., Williamson 2005; Hawthorne and Stanley 2008; Moss 2018 for sympathetic discussion of a norm along these lines.

<sup>20</sup>Of course, the Knowledge Norm of Reliance, much like the Knowledge Norm of Assertion, is controversial. One criticism holds that the Knowledge Norm of Reliance is difficult to square with the idea that rational agents often act on the basis of credences (Schiffer 2007). In response, some have proposed that credences can constitute knowledge (Moss 2018; Beddor and Goldstein 2021). Another line of criticism maintains one can permissibly rely on a false proposition that one justifiably believes to be true (Brown 2008; Neta 2009). In response, proponents of the Knowledge Norm of Reliance often maintain that in such cases an agent violates the norm of reliance, but does so blamelessly—a move that parallels common defenses of KA. See §8 for more discussion.

Together, Promises Warrant Reliance and the Knowledge Norm of Reliance entail:

PROMISES WARRANT KNOWLEDGE A characteristic function of promising to  $\phi$  is to enable the promisee to know that the promisor will  $\phi$ .

This view focuses on the epistemic states of the promisee. However, it has implications for the epistemic states of the promisor. In the vast majority of cases of testimonial knowledge, the speaker transmits knowledge to the hearer. If  $B$  comes to believe  $p$  on the basis of  $A$ 's testimony, but  $A$ 's belief in  $p$  does not amount to knowledge (because it is false, or unjustified, or Gettiered), then usually  $B$ 's belief in  $p$  will not amount to knowledge either.<sup>21</sup> This seems equally true of promising. Suppose  $A$  promises to give  $B$  a shiny toaster. Suppose  $B$  relies on this promise, forming the belief that  $A$  will give them a shiny toaster. If  $A$  didn't actually know that they would give  $B$  a shiny toaster (perhaps  $A$  had no intention of doing so, or  $A$  didn't know whether they could afford one), then it is hard to see how  $B$ 's resulting belief could amount to knowledge. This gives us another argument for KP: in order for promises to fulfill one of their characteristic functions, the promisor will typically need to know that the promise will be fulfilled.

## 7 Promising Involves Asserting

I have argued that there are close parallels between the normative constraints on assertion and the normative constraints on promising. But why do these parallels hold?

Here is a simple answer:

PROMISING INVOLVES ASSERTING Whenever someone promises to  $\phi$ , they thereby assert that they will  $\phi$ .<sup>22</sup>

If promising involves asserting, KP is a direct consequence of KA. No surprise, then, that the normative constraints on promising so closely track the normative constraints on asserting. And no surprise that the functional profile of the former mirrors the functional profile of the latter, at least insofar as both aim at generating knowledge in one's addressee.

Let me briefly highlight some further advantages of Promising Involves Asserting. First, it allows us to simplify our explanation of Moorean absurdities. In order to explain the absurdity of Moorean promises, we needed both KP and KA, since the second conjunct of a Moorean promise is an assertion, rather than a promise. If promising involves asserting, there is no need to invoke two separate norms. Every Moorean promise violates KA.

The second argument is independent of our Moorean data. The following inference pattern seems valid:

---

<sup>21</sup>For discussion, see [Burge 1993](#); [Owens 2006b](#); [Williamson 2000](#); [Kelp 2018](#), among others. For some exceptions to this generalization, see [Lackey 2007](#).

<sup>22</sup>A number of authors have discussed the parallels between promising and asserting, though usually without endorsing anything as strong as the view being put forward here (e.g., [Austin 1946](#); [Scanlon 1990](#); [van Roojen 2020](#)). Though see [Árdall 1968](#); [Thomson 1990](#), both of whom endorse versions of Promising Involving Asserting. For a rejection of this thesis, see [Watson 2004](#).

- (23) a.  $A$  promised to  $\phi$ .  $\Rightarrow$   
 b.  $A$  said that  $A$  will  $\phi$ .

For example, if your neighbor promises to walk your dog, then it seems perfectly natural to report this by saying, “My neighbor said that they will walk my dog.” Just consider how strange it would be for your neighbor to deny that they said any such thing:

- (24) a. *Your neighbor*: I promise to walk your dog while you’re gone.  
 b. *You*: Great! So when you walk him, please make sure to use a harness—  
 c. *Your neighbor*: # Whoa, whoa! I didn’t say that I will walk your dog. I only promised to walk him.  
 d. *You*: Umm... what??

So the inference in (23) certainly seems valid. But, given that the verb “said” is usually reserved to report *assertions*, it can start to look mysterious as to why this inference holds. If promising involves asserting, the mystery is dispelled. Your neighbor did assert that they would walk your dog; they did so by virtue of promising to do so.

I have laid out some arguments in support of the view the promising entails asserting. Does asserting entail promising? This seems implausible. If you ask me about the weather and I tell you it is raining outside, I am not naturally described as having promised that it is raining. What then distinguishes promises from (mere) assertions? Without purporting to offer a full-fledged analysis of the distinction, we can note some hallmarks of paradigmatic promises, drawing on Scanlon 1990. First, paradigmatic promises concern some action that the promisor proposes to perform (or forego performing) in the future. Second, the promisee has an interest in being assured about whether the promisor will perform (or forego performing) the relevant action, or at least the promisor believes the promisee has such an interest. And finally, the promisor makes the promise with the goal of providing that assurance. My assertion that it is raining outside lacks these features.<sup>23</sup>

Taking stock: I have argued that promises are not just *similar* to assertions; rather, promising entails asserting. However, we do not need to take on board this strong conclusion for the main argument in my paper to succeed. My central thesis is just that there is a close parallel between the normative constraints on promising and those on assertion. The claim that promising involves asserting is put forward as a promising explanation for why this parallel holds.

## 8 Is KP Too Demanding?

Let me tackle what I suspect will strike many as the most obvious objection to KP: it is too demanding. This objection comes in a few different forms, which I will consider in turn.

---

<sup>23</sup>I have formulated these as features of “paradigmatic promises” so as to remain noncommittal on the question of whether superficially similar speech acts that lack these features are not promises at all, or rather just atypical promises.

## 8.1 Skepticism About the Future

Most promises concern the future. According to some philosophers of a skeptical persuasion, knowledge of the future is seldom, if ever, attained. By KP, it would seem to follow that future-directed promises are seldom, if ever, permissible.

However, I think we should reject this skeptical outlook. Common sense holds that we can know much about the future. I know that the sun will rise tomorrow; I know that a year from now my work will be less widely read than Plato's. Moreover, most theories of knowledge agree with common sense on this point. As a toy example, take the view that knowledge is justified safe belief, where a belief is safe if it is true in all nearby worlds where it is held on the same basis. This view allows that many beliefs about the future can amount to knowledge. Suppose I promise to attend your graduation next May. If I justifiably believe that I will attend your graduation next May, and this belief is true in all nearby worlds where it is held on the same basis, then the safety analysis predicts my belief qualifies as knowledge. The point generalizes: as long we reject skepticism about the future, KP allows that we can permissibly make future-directed promises.<sup>24</sup>

## 8.2 Excused Promise-Breaking

Another version of demandingness objection comes from cases where someone breaks their promise through no fault of their own. Consider:

**Recital** I promise to attend your recital, and have every intention of doing so. But after making my promise I am regrettably kidnapped, causing me to miss your performance.

Since knowledge is factive (one can only know  $p$  if  $p$  is true), my promise violated KP. Still, some might maintain, there is some sense in which my promise was appropriate.

This is an analogue of a well-known challenge for the Knowledge Norm of Assertion. According to some critics, KA delivers the wrong results in cases of justified false assertions (e.g., Douven 2006: 476-477; Lackey 2007: 603). If I have good evidence that your recital starts at 7pm, it seems appropriate to assert it will start then, even if, unbeknownst to me, the start time is delayed. Defenders of KA have a standard response to this objection, which is to distinguish between norm-compliance and blamelessness. In general when someone violates a norm  $N$ , we tend to excuse their behavior if they reasonably believed they were complying with  $N$ . Suppose your speedometer unpredictably malfunctions,

---

<sup>24</sup>This point helps address a related concern. A referee raises the worry that KP implies that we cannot know whether someone complied with the norm on promising until the time at which the promised action was supposed to occur. For example, KP might seem to imply that when I promise to attend your graduation next May, we cannot assess whether my promise was permissible until next May rolls around. However, if we can know things about the future, presumably we can also know that we know things about the future. So there is no principled reason why I cannot—in the right circumstances—know that I know I will attend your graduation next year. If I can know that I know this, so can others. So KP does not imply that we must always wait until the time at which the promised action was supposed to occur to know that a promise complied with the norm.

leading you to slightly exceed the speed limit. You violate the traffic norms, but this violation was excusable, since you reasonably believed you were complying with those norms. Similarly, when I (falsely) assert that the recital starts at 7, I have violated the norm for assertion. But I was blameless, since I reasonably believed that I knew the recital started then.<sup>25</sup> What goes for asserting goes for promising. My promise to attend your recital violated the norm of promising. But I was blameless, since I reasonably believed I knew that I would attend.

Some might balk at the suggestion that Recital involves an excused norm violation.<sup>26</sup> Adjust the case: I escape my kidnappers, but as I drive to your recital, I see an injured motorist who needs help. If I help the motorist and consequently miss the recital, it seems wrong to say that I am merely *excused* for violating the norm of promising: I *ought* to help the motorist.

In response, we should emphasize that KP is not being put forward as a moral norm. Rather, it is both a conversational norm and an epistemic norm (§4). This is not to say that there is no connection between KP and morality. Indeed, I will argue (§10) that KP helps explain why there is a defeasible moral norm against breaking one's promises; in other words, the conversational/epistemic norm generates a moral norm. For present purposes, the important point is that this derivative moral norm is defeasible: it isn't always morally wrong (all things considered) to violate KP. Consequently, when someone promises to  $\phi$  without knowing that they will  $\phi$ , their speech act is in some sense defective, insofar as it violates (perhaps excusably!) a norm governing promising. But it is not necessarily *morally* defective.

For those who remain unconvinced that Recital involves an excused norm violation, a more concessive reply is also available. Here too, we can take our cue from the norms of assertion literature. Faced with the objection from justified false assertions, defenders of KA typically appeal to the excuse maneuver, as described above. But another option is to modify KA (e.g., [Smithies 2012](#)). For example, one might hold that the norm of assertion is reasonably believing that one knows:

REASONABLE BELIEF KNOWLEDGE NORM OF ASSERTION One ought only assert  $p$  if one reasonably believes one knows  $p$ .<sup>27</sup>

Those attracted to this position could embrace a similar variant of KP:

REASONABLE BELIEF KNOWLEDGE NORM OF PROMISING (RBKP) One ought not promise to  $\phi$  unless one reasonably believes one knows one will  $\phi$ .<sup>28</sup>

<sup>25</sup>Versions of this response are developed in [Williamson 2000](#) and [DeRose 2002, 2009](#). For further discussion of the relation between norm compliance and blameworthiness, see [Benton 2013](#); [Kelp and Simion 2017](#); [Greco 2021](#). For empirical evidence that people often fail to distinguish between complying with a norm and blamelessly violating that norm, see [Turri 2019](#).

<sup>26</sup>Thanks to a referee for helpful questions here.

<sup>27</sup>See [Smithies 2012](#) for a similar proposal (Smithies' 'JK Rule'). See also [Neta 2009](#) for a similar alternative to the Knowledge Norm of Reliance, according to which  $A$  is warranted in relying on  $p$  iff  $A$  justifiably believes that  $A$  knows  $p$ .

<sup>28</sup>We noted in §3 that some might want to explain the infelicity of (18) by appealing to a stronger norm of

According to RBKP, your promise in Recital was permissible: at the time of promise-making, you reasonably believed that you knew you would attend. At the same time, RBKP is faithful to the spirit of KP. It posits a close normative connection between promising and knowledge. And it explains the main data motivating KP. When your neighbor says (1) (“I promise to walk your dog, but I don’t whether I will”), either they reasonably believe they know they will walk your dog or they do not. If they do not reasonably believe they know they will walk your dog, they have violated RBKP. If they do reasonably believe they know they will walk your dog, then they do not reasonably believe they know second conjunct, given the assumption that one cannot reasonably believe obviously inconsistent propositions. RBKP is similarly well-positioned to explain why one can abstain from making a promise by confessing that one doesn’t know whether one will fulfill it, and why it is infelicitous to promise to lose the lottery: one cannot reasonably believe one knows that one’s ticket will lose.<sup>29</sup>

So there are two attractive strategies for handling the objection from excused promise-breaking. The less concessive strategy holds that in such cases the agent violates the norm of promising but does so blamelessly, and moreover that violating this norm is not necessarily a *moral* failing. The more concessive strategy replaces KP with a norm that still preserves a close link between promising and knowledge, but which allows that in such cases the agent’s behavior complied with the norm. RBKP offers one candidate for a such a norm; in §8.3 I’ll mention another. For the purposes of this paper, I want to remain neutral on which strategy we should prefer. The answer to this question will hinge on one’s more general views about the relation between norm violation and blamelessness, and the merits of externalist vs. internalist norms.

### 8.3 Knowing that You Don’t Know

Let me turn to another way of pressing the demandingness objection. According to [Marušić \(2013, 2015\)](#) and [Lieberman \(2019\)](#), sometimes one can appropriately promise to  $\phi$  even though one knows that one does not know that one will  $\phi$ . Let’s start with a case adapted from [Marušić 2015](#):

**Quitting Quandary** You’re a smoker. Your spouse implores you to quit. You want to do so, but you know all too well that the flesh is weak. Indeed, you’ve just read a medical study indicating that most smokers who try to quit fail on their initial attempts. You have no reason to think that you are more likely to succeed where others have failed.

---

promising, according to which one not promising to  $\phi$  unless one knows with certainty that one will  $\phi$ . We could also explore a non-factive analogue of this norm in the spirit of RBKP: one ought not promise to  $\phi$  unless one reasonably believes one knows with certainty that one will  $\phi$ .

<sup>29</sup>RBKP also offers an alternative response to the concern that KP implies we cannot know whether someone complied with the norm of promising until the time at which the promised action was supposed to occur (fn.24). RBKP faces no such worry: even when we are not in a position to know whether *A* knows they will fulfill their promise, we might still be in a position to know whether *A* reasonably believes they know they will fulfill their promise.



Marušić contends that it is rational to promise to quit smoking. But in this scenario, you know that you don't know whether you will quit.

However, our arguments provide reason to resist this counterexample.<sup>30</sup> As an initial observation, note that in Quitting Quandary it still sounds absurd to conjoin your promise with an acknowledgment that you don't know whether you will carry through:

(25) # I promise to quit smoking, but  $\left\{ \begin{array}{l} \text{I don't whether I will} \\ \text{I might not quit} \end{array} \right\}$ .<sup>31</sup>

Moreover, we should ask why it seems rational to promise to quit smoking in this scenario. The obvious answer is that quitting would be beneficial to both you and your spouse. By promising to quit, you give yourself an additional incentive to do so, making it somewhat more likely that you will quit. Arguably, then, you have a prudential—and perhaps moral—reason to promise to quit.<sup>32</sup> However, we can come up with analogous cases involving assertion. Suppose someone offers you a million dollars to assert that it will rain tomorrow, even though you have no idea what the weather will be like. Their offer provides a practical reason to assert something that you do not know. But this does not mean that KA is false, or even that KA is defeasible. After all, it still sounds terrible to assert the Moorean conjunction, “It will rain but I don't know it will”, suggesting that KA is still in force. Rather, all this shows is that sometimes one has a compelling reason to violate KA.

This diagnosis offers a natural way of reconciling intuitions about Quitting Quandary with KP. In Quitting Quandary, you have a practical reason to make a promise while knowing that you don't know whether you will fulfill it. But this does not show that KP is false or defeasible. The fact that (25) sounds terrible in this context provides strong evidence that your promise is in some sense defective, even if you have a compelling reason to make this defective promise.

Some might think that even if this response handles Quitting Quandary, further trouble is in store. According to [Lieberman 2019](#), the source of an agent's uncertainty about whether they will fulfill their promise makes an important normative difference. On her view, if the reason that you don't know whether you will fulfill your promise is that you harbor serious doubts your own capacities or strength of will (what she calls “internal uncertainty”), then your promise is impermissible. But if the reason that you don't know

---

<sup>30</sup>Not everyone shares Marušić's judgments about cases like Quitting Quandary. For example, [Brinkerhoff 2021](#) argues that either you must have some special epistemic reason to think you are more likely to succeed where others have failed or it would be irresponsible to promise to quit. In a similar vein, [D'Cruz and Kalef 2015](#) argue that in this sort of example, responsible agents will promise to *try* to quit, rather promising to quit.

<sup>31</sup>To my knowledge, Marušić never explicitly offers a diagnosis what is wrong with sentences like (25). However, [Marušić 2013, 2015](#) does say that sincerely promising to do something requires believing that one will do it. It's just in cases like Quitting Quandary, one's belief is not made rational by the available evidence, but rather by the practical reasons in favor of quitting smoking. However, we'll see later (§9.2) that a belief requirement on promising faces a dilemma: either it is too weak to explain the data, or it fails to be a genuine alternative to KP (or RBKP).

<sup>32</sup>Emphasis on “arguably”. After all, if you promise to quit smoking and then fail to do so, you risk having misled your spouse. And this may in turn provide a reason to refrain from promising. See [Brinkerhoff 2021](#) for a related criticism of Marušić's view.

whether you will fulfill your promise is that you don't know whether external circumstances will thwart your fulfillment ("external uncertainty"), then your promise might still be permissible.<sup>33</sup> Quitting Quandary is most naturally understood as a case of internal uncertainty. Perhaps, then, cases of external uncertainty present stronger counterexamples to KP.

To evaluate this suggestion, let us consider a case of external uncertainty:

**Travel Plans** Jesse is visiting his family for the holidays. His flight is scheduled to arrive at 5pm. Unfortunately, he's flying JetBlue, and he recently read that JetBlue's flights are frequently delayed. Jesse has no special reason to think his flight will avoid this fate. He's on the phone with his mother, who wants to know his arrival time so she can plan accordingly.

Is it appropriate for Jesse to promise to arrive by 5pm? Intuitions may differ, but it seems to me the answer is *no*. Doing so would make a stronger commitment than Jesse can reasonably expect to fulfill. Instead, he should undertake a more hedged commitment—e.g., "I promise to do my best to arrive by 5pm", or "I promise to take a flight scheduled to arrive at 5pm, but I can't guarantee it will arrive on time."

To draw this intuition out, consider how Jesse's mother would react if he conjoined an unqualified promise with an admission of his lack of knowledge:

(26) # I promise to arrive by 5pm, but I don't know that I will—there's a decent chance my flight will be delayed.

To my ears this sounds bad—just as infelicitous as any other Moorean promise.

We can reinforce this judgment on more theoretical grounds. I've argued that someone who is in a position to promise to  $\phi$  is also in a position to assert that they will  $\phi$  (Promise-Assertion Bridge). So if Jesse were in a position to promise to arrive by 5pm, he should also be in a position to assert he will arrive by then. But consider the corresponding assertion:

(27) # I will arrive by 5pm, but I don't know that I will—there's a decent chance my flight will be delayed.

This sounds just as absurd any other Moorean assertion.

What if Jesse had no reason to expect a delay? The more remote and unforeseeable the possibility of delay, the more inclined we are to think it would be appropriate for Jesse to promise to arrive by 5pm. But this observation fits comfortably with my view. The more remote and unforeseeable the possibility of delay, the more reasonable it would be for Jesse to believe that he knows he will get in by 5pm. And so we can explain the appropriateness of his promise either by saying that he was fully excused for violating the norm of promising (since he reasonably believed he was complying with the norm), or else by saying that he fulfilled the norm of promising altogether (if we replace KP with RBKP).

---

<sup>33</sup>See [Lieberman 2021](#) for a similar distinction applied to wedding vows.

For readers who do not share my judgments, all is not lost. Another way of modifying KP is to require that you know you will fulfill your promises *unless you are excused from doing so*:

KP (WEAKENED) One ought not promise to  $\phi$  if one does not know (one will  $\phi$  unless one is excused from fulfilling one's promise).

KP (Weakened) allows that Jesse could promise to arrive by 5pm (assuming a flight delay qualifies as a valid excuse). My own inclination is to regard KP (Weakened) as *too* weak, since to my ears (26) sounds defective. But my goal here is not to dictate readers' intuitions; for those who disagree, KP (Weakened) still captures much of the spirit and motivations behind KP.<sup>34</sup>

Where does this leave us? Many will regard a knowledge norm on promising as overly demanding. This section sought to defuse this concern. I considered a few ways of fleshing out the "demandingness" objection, and argued some version of a knowledge norm can withstand them all.

## 9 Alternative Explanations?

I've motivated a knowledge norm on promising through inference to the best explanation: it explains a variety of linguistic data (§§4-5), as well as a characteristic function of promises (§6). Still, some might wonder: is a knowledge norm really needed? Could some other constraint on promising explain the data equally well? In this section, I'll examine some of the most promising alternative constraints. I'll argue that each fails to provide a satisfactory alternative explanation of the data.

### 9.1 Promising and Intending

According to many philosophers, there is a close connection between promising and intending. Here's one way of fleshing out this connection:

INTENTION NORM OF PROMISING (IP) One ought not promise to  $\phi$  if one does not intend to  $\phi$ .

---

<sup>34</sup>The choice between KP, RBKP, and KP (Weakened) is also relevant when an agent wants to make a promise, but thinks that the promisee is likely to refuse to accept the promise, or else is likely to release them from their promise. Suppose you want to offer to pick your friend up from the airport. Alas, you have every reason to doubt that your friend will take you up on it: you know she wouldn't want to risk inconveniencing you, and she has expressed concerns about your recent spate of driving accidents. According to a flat-footed application of KP and RBKP, you shouldn't promise to pick them up from the airport; rather, you should make a conditional promise—e.g., "I promise to pick you up from the airport if you'll let me." By contrast, KP (Weakened) allows that you can straightforwardly promise to pick them up from the airport, since if they decline your promise or release you from it, you are excused from performance. (Even here, however, proponents of KP and RBKP have some options for mimicking the predictions of KP (Weakened). Perhaps when you promise to pick your friend up from the airport, your promise is implicitly conditional on the promisee's uptake. And you do know (or reasonably believe you know) that you will fulfill your promise if this background condition is satisfied.)

A notable feature of intentions is that they close off possibilities. This idea is emphasized by [Bratman 1987](#), who takes intention to be a *settling attitude*. To illustrate with Bratman's example, if I merely desire to spend the afternoon at the library, then I may not have made up my mind whether I will do so. Consequently, I might find myself deliberating about whether to head to the library or go to the movies instead. By contrast, if I *intend* to spend the afternoon at the library, I have in some sense made up my mind to do so: I have "closed off" the possibility that I will spend the afternoon at the movies instead ([1987: 18-19](#)). Some might wonder: could this settling feature of intention explain our data?

Only if we explain settling in terms of knowledge:

KNOWLEDGE NORM OF INTENTION One ought not intend to  $\phi$  unless one knows (or at least reasonably believes one knows) one will  $\phi$ .

Given a Knowledge Norm of Intention, IP offers a simple explanation of our Moorean data. By IP, your neighbor should utter the first conjunct of (1) ("I promise to walk your dog") only if they intend to walk your dog. By a Knowledge Norm of Intention, they should have this intention only if they know (or reasonably believe they know) that they will walk your dog. But then they cannot know (or reasonably believe they know) the second conjunct.<sup>35</sup>

However, there are two problems for this proposal. The first is that the Knowledge Norm of Intention is too strong: we can rationally intend to do things in full awareness that we don't know whether we will succeed. Suppose we are co-authoring a paper. I ask when you will finish your share of the revisions. You reply:

- (28) ✓ I intend to finish my revisions by Friday, but I don't know whether I'll manage to do so—I've got a lot on my plate.

Your reply is free from any whiff of Moorean absurdity.

But swap out "intend" with "promise" and your utterance crashes:

- (29) # I promise to finish my revisions by Friday, but I don't know whether I'll manage to do so—I've got a lot on my plate.

The proposal under consideration does not explain this difference in felicity.

On its own, this objection may not be insurmountable. Perhaps some conative states are stronger than intention. For example, [Lieberman 2020](#) proposes that *resolutions* are mental states that involve a firmer commitment than mere intentions.<sup>36</sup> And it does sound rather odd to say:

<sup>35</sup>Some might think that we can explain the Moorean data with a slightly weaker requirement: namely, that you shouldn't intend to  $\phi$  if you believe that you don't know whether you will  $\phi$ . However, this weaker requirement would not explain the conversational data involving presuppositions. Recall (21), where Red promises to be home by Christmas, and Kitty asks how he knows he'll be home by then, presupposing that Red does know this. The weaker requirement does not explain why Kitty's presupposition is warranted. After all, Red could fail to believe that he doesn't know he will be home by Christmas without in fact knowing he will be home by then ( $\neg(B\neg Kp) \nRightarrow Kp$ ).

<sup>36</sup>See also [Fruh 2019](#) for related discussion.

(30) ? I resolve to finish my revisions by Friday, but I don't know whether I will.

Perhaps, then, we should follow [Lieberman 2020](#) and hold that promises express *resolutions*, and combine this proposal with a knowledge norm on resolutions.

But even if the first problem can be overcome in this fashion, a more important issue looms: the resulting view is not a genuine *alternative* to KP. After all, if one should only promise to  $\phi$  if one intends (or resolves) to  $\phi$ , and if one should only intend (or resolve to)  $\phi$  if one knows one will  $\phi$ , it follows that one shouldn't make a promise unless one knows one will keep it. But this is exactly what KP says! So the relevant package of principles is not weaker than KP. Rather, it is just one way of implementing a knowledge norm of promising.

## 9.2 Promising and Believing

Analogous problems arise if we try to explain the data in terms of a connection between promising and believing, e.g:

BELIEF NORM OF PROMISING (BP) One ought not promise to  $\phi$  if one does not believe one will  $\phi$ .

BP will only explain the data if knowledge is the norm of belief:

KNOWLEDGE NORM OF BELIEF One ought not believe  $p$  if one does not know (or reasonably believe one knows)  $p$ .

Here too, we might worry whether the Knowledge Norm of Belief is too strong, since there seems to be nothing irrational about claiming:

(31) ✓ I believe I'll finish my revisions by Friday, but I don't know if I will—I've got a lot on my plate.<sup>37</sup>

One might try to get around this point by suggesting that promising requires a stronger cognitive state than belief, a state such as *being convinced* or *being certain*. Perhaps *this* state is subject to a knowledge norm.<sup>38</sup> But even if we grant this point, a more serious difficulty lies in wait: combining BP with a Knowledge Norm of Belief *entails* KP. Once again we have not uncovered a rival to KP. We have just found another path to the same destination.<sup>39</sup>

---

<sup>37</sup>See [Hawthorne et al. 2016](#); [Beddor and Goldstein 2018](#).

<sup>38</sup>For relevant discussion, see [Beddor 2020a](#); [Goodman and Holguín forthcoming](#).

<sup>39</sup>Similar points apply if we try to explain our data using a control constraint on promising, according to which you should only promise to that which is under your control, or that which you rationally believe to be under your control ([Lieberman 2021](#)). A control requirement on promising would only explain the data if there is a knowledge condition on control, according to which if  $\phi$ ing is under your control, then you know (or reasonably believe you know) whether you will  $\phi$ . But if we endorse to a knowledge condition on control, the resulting view is not an alternative to KP (or RBKP).

### 9.3 Promising and Posturing

Let me briefly address another alternative norm on promising, inspired by recent work on imperatives by [Mandelkern 2021](#). Mandelkern starts by observing that we find Moorean clashes involving imperatives—what he calls “practical Moore sentences”. Suppose you have just been kidnapped. Outraged, you exclaim:

(32) Release me!

Your order seems perfectly in order. But now imagine conjoining your command with an admission that you don’t know whether your captors will comply:

(33) # Release me! I don’t know whether you will.

This sounds very strange.

It might be tempting to explain this infelicity using a knowledge norm on imperatives: you should only order someone to  $\phi$  if you know they will  $\phi$ . But this seems much too strong. In the scenario just described, it is perfectly appropriate to utter (32) on its own. But you don’t know that your captors will release you. Indeed, you know that you don’t know this.

So why is (33) infelicitous? Mandelkern proposes that when you issue a command, you must adopt a posture or pretense of knowledge:

POSTURING NORM ON IMPERATIVES When ordering someone to  $\phi$ , you must act as if you know that they will  $\phi$ .

In further work, [Mandelkern and Dorst 2022](#) defend a similar posturing norm on assertion: when asserting  $p$ , you must act as if you know  $p$ .<sup>40</sup> A natural thought is that we could extend the posturing norm to encompass promises as well:

POSTURING NORM ON PROMISING (PP) When promising to  $\phi$ , you must act as if you know you will  $\phi$ .

PP offers an alternative story about why Moorean promises are absurd. Acknowledging that you don’t know whether you will  $\phi$  is inconsistent with acting as if you know you will  $\phi$ . So no one could utter a Moorean promise while abiding by PP.

However, PP struggles to explain our additional arguments for KP. One such argument was that we often abstain from making a promise when we don’t know whether will carry through, as revealed by our foray into *That 70s Show* fan fiction (§5.1). KP explains this tendency; PP does not. Even though Red doesn’t know he will return home from the war, there is nothing to prevent him from acting as if he knows this. For much the same reason, PP does not explain the infelicity of lottery promises (§5.3): nothing is stopping you from *pretending* to know that your ticket will lose.

Or take the argument from conversational presuppositions (§5.2). We saw that promising to  $\phi$  typically licenses a presupposition that the promisor knows they will  $\phi$ , as illustrated by our variant of the Red-Kitty dialogue:

---

<sup>40</sup>See [Dinges 2024](#) for a related proposal.

- (21) a. *Red*: I promise I'll be home by Christmas.  
 b. *Kitty*: How do you know you'll be back by then?  
 c. *Red*: # Sheesh, I didn't say I *knew* I would be back by then. I just promised to be back by then!

KP explains why this reply is so strange. PP does not. PP merely requires that one pretend to know that one will comply with one's promise at the time of promising. It does not impose the further requirement that one keep up the pretense at all subsequent points in the conversation.

The comparison with imperatives is revealing. Consider:

- (34) a. *Professor*: Turn in your papers by 5pm!  
 b. *Student*: # How do you know I'll turn my paper in by then?  
 c. *Professor*: Sheesh, I didn't say I *knew* you would turn it in by then. I just told you to do it!

Here, the knowledge-presupposing question is infelicitous. The knowledge-denying reply is perfectly in order.

This contrast suggests a more general lesson. At first blush, it is tempting to try to provide a unified explanation of Moorean promises and practical Moore sentences. But, on closer examination, there are important differences between the data involving promising and those involving imperatives. An adequate theory of speech acts should explain these differences. Consequently, we should not look for a one-size-fits-all norm of promising and commanding.<sup>41</sup>

For our purposes, we can remain neutral on the correct explanation of practical Moore sentences. Perhaps a posturing norm of imperatives holds the answer; perhaps not. The important point is that a posturing norm on promises is too weak to explain the full range of data. We have yet to find a rival to KP that does the trick.

This completes my development and defense of a knowledge norm of promising. I now turn to consider its normative payoff.

## 10 The Source of Promissory Obligation

Promises give rise to obligations. If I promise to water your plants, I incur an obligation to do so. But how does this work, exactly? This question is at the heart of a central—perhaps

<sup>41</sup>Cf. Ninan 2005, who also seeks a unified explanation of Moorean imperatives and Moorean promises. Ninan's Moorean data are importantly different from our examples. Ninan focuses on Moorean clashes of the form, "Shut the door! You are not going to shut the door" and "I promise to mop the floor, but I'm not going to." Ninan's solution is that imperatives and promises are subject to a weak belief constraint. With imperatives, the constraint is that you shouldn't order someone to  $\phi$  if you believe that they won't  $\phi$ . With promises, the constraint is that you shouldn't promise to  $\phi$  if you believe you will not  $\phi$ . But even if this constraint explains Ninan's data, it is too weak to explain the Moorean data that has been our focus. Moreover, the observations in this section call into question whether we should want a unified treatment of Moorean imperatives and Moorean promises in the first place.

*the* central—debate in the literature on promising. How does promising to  $\phi$  generate an obligation to  $\phi$ ?

One popular family of answers comes from expectation theories. While there are different ways of developing expectation theories, they all agree that when I promise to water your plants, I give you reason to expect that I will water them.<sup>42</sup> So if I fail to hydrate them, I have misled you. But we have a defeasible obligation not to voluntarily mislead others, which is why promises give rise to obligations.

This is a very attractive diagnosis. But it faces a major challenge. Why does my promise to water your plants give you reason to expect I will carry through? What is it about promising that justifies this expectation? Some have worried that this cannot be answered without circularity. According to this worry, the reason why my promise gives you reason to expect I will carry through is that you expect I will fulfill my promissory obligations. But this takes for granted the very thing we were trying to explain.<sup>43</sup>

A knowledge norm of promising offers a new, simple solution to this long-standing problem. Earlier we invoked the plausible principle that we expect others to abide by conversational norms (§5.3):

EXPECTED COMPLIANCE When someone performs a particular speech act, we are defeasibly entitled to expect that they satisfy the norms governing the speech act.

We already encountered evidence in favor of this principle, courtesy of our argument from conversational presuppositions. Take assertion first. When someone asserts  $p$ , it can be appropriate to ask, “How do you know that?”, which presupposes the speaker knows  $p$  (as illustrated by (20)). As we saw, a number of philosophers have taken these data to support KA. But these data only support KA given Expected Compliance: it is only because we are defeasibly entitled to expect that speakers satisfy the norm of assertion (KA) that we consider it appropriate to ask questions that presuppose they satisfy this norm.

This point carries over to promising. We observed that when someone promises to  $\phi$ , it is appropriate to ask how they know that they will  $\phi$ —a question that presupposes the promisor has this knowledge (as illustrated by (21)). KP explains why. But it only does so given Expected Compliance: it is only because we are defeasibly entitled to expect that speakers satisfy the norm of promising (KP) that we consider it perfectly appropriate to ask questions that presuppose they satisfy this norm.<sup>44</sup>

<sup>42</sup>Expectation theories include reliance views (which we encountered in §6), as well as assurance views (Scanlon 1990, 1998) and trust views (Friedrich and Southwood 2011).

<sup>43</sup>See e.g. Kolodny and Wallace 2003 for arguments that Scanlon’s version of an expectation view faces a circularity problem.

<sup>44</sup>As noted in fn. 13, a version of Expected Compliance is also presupposed by Gricean pragmatics. Still, one might wonder: why does Expected Compliance hold? For our purposes, we do not need to commit to a particular answer; it is enough that it *does* hold. That said, let me mention one possibility. Simion and Willard-Kyle 2023 defend the much more general principle that whenever there is some operative norm obligating an agent  $A$  to  $\phi$ , we have a default entitlement to trust that  $A$  will  $\phi$ . As they observe, many norms are routinely followed: drivers generally drive on the correct side of the road and refrain from running lights; shoppers generally avoid cutting in supermarket queues. Simion and Willard-Kyle also point to a rich body of empirical evidence that people have a strong tendency to abide by social norms. Perhaps, then, Expected Compliance is just a special case of this much more general principle.



Combining KP and Expected Compliance yields an elegant solution to the circularity problem. When I promise to water your plants, you are defeasibly entitled to expect that I will satisfy the norms governing my promise. Consequently, you are defeasibly entitled to expect that I know I will water your plants (by KP). Now, knowledge is factive: I cannot *know* that I will water your plants unless I will in fact do so. Therefore, you are defeasibly entitled to expect that I will water your plants. The circularity problem is solved. And it is solved using just two ingredients, Expected Compliance and KP, both of which are supported by independent data.<sup>45,46</sup>

Some might worry that my solution is unable to distinguish between mere assertions and promises with respect to how they invite reliance. Consider a pair of cases:

**League Assertion** We are discussing our plans for the weekend. I tell you that I'll be at the little league game tomorrow, since I've volunteered to umpire.

**League Promise** We are discussing our plans for the weekend. I *promise* to be at the little league game tomorrow, since I've volunteered to umpire.

If I don't show up at the little league game, you have more of a complaint against me in the second case than in the first. But if knowledge is the norm of both assertion and promising, then by Expected Compliance you have a defeasible reason to expect I know that I will be at the game in both cases. So what explains the normative difference between the two scenarios?<sup>47</sup>

In response, it will be helpful to ask, why did I *promise* to be at the game in League Promise? On the most natural way of filling in the details, I made a promise because I thought it was important to you that I would be at the game, or at least that I thought it was important to you to know where I would be tomorrow. (If we were merely making

---

<sup>45</sup>What if we replace KP with RBKP? Since RBKP is not a factive norm, the solution does not go through quite the same. But a closely related solution applies. By RBKP and Expected Compliance, you are entitled to expect that I reasonably believe I know I will water your plants. Since you take my belief to be reasonable, you are entitled to expect I have good reason to believe I will water your plants. This expectation, in turn, gives you some reason to think I will water your plants. By contrast, no such solution works if we adopt the posturing norm (PP). By PP and Expected Compliance, you would only be reasonably entitled to expect that I am acting as if I know I will water your plants. But this does not entitle you to believe that I will water them.

<sup>46</sup>Another proposed solution to the circularity problem, due to [Kolodny and Wallace 2003](#), combines a practice-based theory of promissory obligation with an expectation-based theory. According to their hybrid view, there is an initial practice-based moral obligation to fulfill one's promises—an obligation that is not grounded in the promisee's expectations. But, given that promisees are aware of this practice-based moral obligation, they are defeasibly entitled to expect that promisors will fulfill their promises, which generates a secondary, expectation-based moral obligation to fulfill one's promises. However, some philosophers doubt whether practice-based views underwrite a sufficiently robust moral obligation to keep one's promises; indeed, this skepticism was part of what motivated the development of expectation-based theories (cf. [Scanlon 1990](#)). This raises the worry that the hybrid theory inherits the problems with practice-based accounts—problems that expectation-based theories were designed to avoid. (See e.g., [Mason 2005](#).) My proposal avoids this concern. It does not posit any practice-based *moral* obligations. Rather, it grounds the promisee's expectation of compliance in a general expectation that speakers will follow the norms governing their speech acts, together a non-moral norm of promising (KP).

<sup>47</sup>Thanks to a referee for raising this question.

small talk and I had no reason to suspect that my whereabouts on the morrow mattered to you, it would be distinctly odd to make a *promise*, rather than an offhand assertion.) This is what we should expect, given our earlier remarks about what distinguishes mere assertions from paradigmatic promises. As we noted in §6, a hallmark of paradigmatic promises is that the promisee has an interest in being assured of whether the promisor will perform the promised action (or at least the promisor believes they have such an interest), and the promisor makes the promise with the goal of providing that assurance. By contrast, in League Assertion, there is no reason to think that my presence at the game matters to you.

I submit that this explains the normative difference between the two cases. More generally, the following seems highly plausible:

FORESEEN RELIANCE-OBLIGATION CONNECTION *Ceteris paribus*, it is worse for *A* to intentionally cause *B* to falsely believe *p* if *A* knows (or believes) that *B* has an interest in forming a true belief about *p*, and that *B* is likely to rely on their belief about *p* in their future actions, than if *A* lacks such knowledge (beliefs).

As evidence that a principle along these lines is driving our intuitions, suppose we tweak the League Assertion to control for these features:

**League Assertion Variant** It's important to you that we meet up tomorrow, and you communicate this importance to me. I assert, "I'll be at the little league game tomorrow."

In this version of the case, it seems you have a genuine complaint against me if I don't show up. You might reasonably protest, "But I was counting on you, and you told me that you would be there!" Indeed, one might go a step further: one might think that once we make these alterations to the case, my assertion that I will be at the game could be reasonably construed as a promise. After all, one does not have to utter the magic words "I promise" for a promise to be made. In many contexts, a bare assertion of what one will do suffices. Arguably, League Assertion Variant is such a context. But even if one doesn't agree with this further claim, the main point still holds: Foreseen Reliance-Obligation Connection offers a plausible explanation for the normative difference between League Assertion and League Promise.

Foreseen Reliance-Obligation Connection is not a consequence of the core framework presented here (KA, KP, or Expected Compliance). But it is compatible with this framework, and it is independently plausible. Moreover, expectation theories are often formulated in ways that encode some version of this idea. For example, Scanlon's official formulation of his expectation theory (Scanlon 1990) is not a blanket prohibition on voluntarily misleading others. Rather, it prohibits voluntarily misleading another party about your future actions *when you know that they have an interest in being assured about your future actions*.<sup>48</sup>

<sup>48</sup>See Scanlon's 'Principle F' (Scanlon 1990: 208). We also see a form of Foreseen Reliance-Obligation Connection in contract law. Recall that the doctrine of promissory estoppel allows otherwise unenforceable promises to be enforced only if the promisee *foreseeably* relies on the promise to their detriment.

## 11 Conclusion

There is something absurd about making a promise and, in the same breath, admitting that you do not know whether you will carry through. But why? Mainstream views of promising fail to provide a satisfactory explanation. In this paper, I've advanced a new epistemic condition on promising that fares better. According to the view put forward here, promising is subject to a knowledge norm (KP): one should only make a promise if one knows one will keep it. This norm explains the absurdity of Moorean promises, and it is supported by a wealth of auxiliary data.

A knowledge norm of promising has important implications for a number of central debates about promising. It provides reason to think that the norms governing promising are importantly stronger than the norms governing intention and belief. It offers insight into the relationship between promising and asserting. And it delivers a new solution to the circularity problem for expectation-based views of the sources of promissory obligation.<sup>49</sup>

## References

- G.E.M. Anscombe. Rules, rights, and promises. *Midwest Studies in Philosophy*, 3(1):318–323, 1978.
- Páll Árdall. And That's a Promise. *The Philosophical Quarterly*, 18(72):225–237, 1968.
- J.L. Austin. Other Minds. *Proceedings of the Aristotelian Society*, Supplementary Volume 20:148–187, 1946.
- J.L. Austin. *How to do Things With Words*. Clarendon Press, Oxford, 2nd edition, 1962.
- Kent Bach and Robert Harnish. *Linguistic Communication and Speech Acts*. MIT Press, Cambridge MA, 1979.
- Bob Beddor. New Work for Certainty. *Philosophers' Imprint*, 20(8):1–25, 2020a.
- Bob Beddor. Certainty in Action. *The Philosophical Quarterly*, 70(281):711–737, 2020b.
- Bob Beddor and Simon Goldstein. Believing Epistemic Contradictions. *The Review of Symbolic Logic*, 11(1):87–114, 2018.
- Bob Beddor and Simon Goldstein. Mighty Knowledge. *Journal of Philosophy*, 118(5):229–269, 2021.
- Matthew Benton. Two More for the Knowledge Account of Assertion. *Analysis*, 71:684–687, 2011.
- Matthew Benton. Dubious Objections From Iterated Conjunctions. *Philosophical Studies*, 162(2):355–358, 2013.
- Matthew Benton. Knowledge is the norm of assertion. In Sosa, Turri, and Roeber, editors, *Contemporary Debates in Epistemology*. Wiley, third edition, 2023. URL <https://philpapers.org/archive/BENKIT-3.pdf>.
- Michael Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.
- Anna Brinkerhoff. The promising puzzle. *Philosophers' Imprint*, 21(22):1–17, 2021.
- Jessica Brown. Knowledge and Practical Reason. *Philosophy Compass*, 3(6):1135–1152, 2008.
- Tyler Burge. Content Preservation. *The Philosophical Review*, 102(4):457–488, 1993.
- Jason D'Cruz and Justin Kalef. Promising to Try. *Ethics*, 125(3):797–806, 2015.

<sup>49</sup>I am grateful to audiences at UT-Austin, the New York Philosophy of Language Workshop, the NYU Abu Dhabi Reasoning and Normativity Workshop, and the University of Florida for feedback. Special thanks to Kyle Blumberg, Sam Carter, Eliran Haziza, Ben Holguín, Nico Kirk-Giannini, Harvey Lederman, two anonymous referees, and the editors of *Ethics* for helpful comments on an earlier draft.

- Keith DeRose. Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2):167–203, 2002.
- Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Clarendon Press, Oxford, 2009.
- Alexander Dinges. Assertion and certainty. *The Philosophical Quarterly*, 74:169–186, 2024.
- Kevin Dorst. Can the Knowledge Norm Co-Opt the Opt-Out? *Thought*, 3(4):273–282, 2014.
- Igor Douven. Assertion, Knowledge, and Rationality Credibility. *Philosophical Review*, 115(4):449–485, 2006.
- Jane Friedman. Question-Directed Attitudes. *Philosophical Perspectives*, 27(1):145–174, 2013.
- Daniel Friedrich and Nicholas Southwood. Promises and Trust. In Hanoch Sheinman, editor, *Promises and Agreements: Philosophical Essays*. 2011.
- Kyle Fruh. Promising's neglected siblings: Oaths, vows, and promissory obligation. *Pacific Philosophical Quarterly*, 100(3):858–880, 2019.
- Jeremy Goodman and Ben Holguín. Thinking and being sure. *Philosophy and Phenomenological Research*, forthcoming.
- Daniel Greco. Justification and Excuses in Epistemology. *Noûs*, 55(3):517–537, 2021.
- Paul Grice. Logic and Conversation. In *The Logic of Grammar*, pages 64–75. Dickensen, Encino, CA, 1975.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, New York, 2004.
- John Hawthorne and Jason Stanley. Knowledge and Action. *Journal of Philosophy*, 105(10):571–590, 2008.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief Is Weak. *Philosophical Studies*, 173(5):1393–1404, 2016.
- Alexander Heape. Promises, Obligation, and Reliance. *Philosophy and Phenomenological Research*, 104(1):150–170, 2022.
- David Hume. *A Treatise of Human Nature*. Oxford University Press, New York, 1739/1978.
- Christoph Kelp. Assertion: A Function First Account. *Noûs*, 52(2):411–442, 2018.
- Christoph Kelp and Mona Simion. Criticism and Blame in Action and Assertion. *Journal of Philosophy*, 114(2):76–93, 2017.
- Nico Kolodny and Jay Wallace. Promises and Practices Revisited. *Philosophy and Public Affairs*, 31(2):119–154, 2003.
- Jennifer Lackey. Norms of assertion. *Noûs*, 41(4):594–626, 2007.
- Alida Liberman. Permissible Promise-Making Under Uncertainty. *Journal of the American Philosophical Association*, 5:468–486, 2019.
- Alida Liberman. The Mental States First Theory of Promising. *Dialectica*, 74(4):665–692, 2020.
- Alida Liberman. For Better or for Worse: When Are Uncertain Wedding Vows Permissible? *Social Theory and Practice*, 47(4):765–788, 2021.
- Neil MacCormick. Voluntary Obligations and Normative Powers. *Aristotelian Society Supplementary Volume*, 46:59–102, 1972.
- Matthew Mandelkern. Practical Moore Sentences. *Noûs*, 55(1):39–61, 2021.
- Matthew Mandelkern and Kevin Dorst. Assertion is Weak. *Philosophers' Imprint*, 22, 2022.
- Berislav Marušić. Promising Against the Evidence. *Ethics*, 123:297–317, 2013.
- Berislav Marušić. *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford University Press, Oxford, 2015.
- Elinor Mason. We make no promises. *Philosophical Studies*, 123:33–46, 2005.
- G.E. Moore. *Commonplace Book: 1919-1953*. Allen and Unwin, London, 1962.
- Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, Oxford, 2018.
- Ram Neta. Treating Something as a Reason for Action. *Noûs*, 43:684–699, 2009.

- Dilip Ninan. Two puzzles about deontic necessity. In Gajewski, Hacquard, Nickel, and Yalcin, editors, *New Work on Modality*, volume 51. MIT Working Papers in Linguistics, 2005.
- David Owens. A Simple Theory of Promising. *The Philosophical Review*, 115(1):51–77, 2006a.
- David Owens. Testimony and Assertion. *Philosophical Studies*, 130:105–129, 2006b.
- David Owens. *Shaping the Normative Landscape*. Oxford University Press, Oxford, 2014.
- H.A. Prichard. The obligation to keep a promise. In MacAdam, editor, *Moral Writings*, page 257–266. Clarendon Press, Oxford, 2002.
- Joseph Raz. Voluntary Obligations and Normative Powers. *Proceedings of the Aristotelian Society*, (46):79–102, 1972.
- Joseph Raz. Promises and Obligations. In *Law, Morality and Society: Essays in Honour of H.L.A. Hart*, pages 210–228. Clarendon Press, Oxford, 1977.
- Yasha Sapir and Peter van Elswyk. Hedging and the Ignorance Norm on Inquiry. *Synthese*, 199(3-4): 5837–5859, 2021.
- Thomas Scanlon. Promises and Practices. *Philosophy and Public Affairs*, 19(3):199–226, 1990.
- T.M. Scanlon. *What We Owe to Each Other*. Harvard University Press, Cambridge, MA, 1998.
- Jonathan Schaffer. Knowledge in the Image of Assertion. *Philosophical Issues*, 18(1):1–19, 2008.
- Stephen Schiffer. Interest-Relative Invariantism. *Philosophy and Phenomenological Research*, 75 (1):188–195, 2007.
- John Searle. *Speech Acts: An Essay in Philosophy of Language*. Cambridge University Press, Cambridge, 1969.
- Mona Simion and Christopher Willard-Kyle. Trust, Trustworthiness, and Obligation. *Philosophical Psychology*, 2023. doi:<https://doi.org/10.1080/09515089.2023.2223221>.
- M.A. Slote. Assertion and Belief. In Dancy, editor, *Papers on Language and Logic*. Keele University Library, Keele, 1979.
- Declan Smithies. The Normative Role of Knowledge. *Noûs*, 46(2):265–288, 2012.
- Jason Stanley. Knowledge and Certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- Judith Jarvis Thomson. *The Realm of Rights*. Harvard University Press, Cambridge, 1990.
- John Turri. Prompting Challenges. *Analysis*, 70(3), 2010.
- John Turri. The Express Knowledge Account of Assertion. *Australasian Journal of Philosophy*, 89 (1):37–45, 2011.
- John Turri. Excuse Validation: A Cross-Cultural Study. *Cognitive Science*, 43(8), 2019.
- Peter Unger. A Defense of Skepticism. *The Philosophical Review*, 80(2):198–219, 1971.
- Peter Unger. *Ignorance: The Case for Skepticism*. Clarendon Press, Oxford, 1975.
- Mark van Roojen. Promising and Assertion. In Goldberg, editor, *The Oxford Handbook of Assertion*, pages 179–200. Oxford University Press, Oxford, 2020.
- Gary Watson. Asserting and Promising. *Philosophical Studies*, 117(1/2):57–77, 2004.
- Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000.
- Timothy Williamson. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55(219):213–235, 2005.