

The Promising Puzzle

Anna Brinkerhoff

Consider the following three claims.¹

- (1) The Evidence Claim: it is rational for you to believe that p only if p is supported by your evidence.
- (2) The Sincerity Claim: your promise to φ is sincere only if you believe that you will φ given that you promise to φ .²
- (3) The Permissibility Claim: it is permissible for you to promise to φ only if your promise is sincere and the belief required for sincerity is rational.

On the face of it, each of these claims is plausible. Here's what's interesting: if all three are true, it seems that we shouldn't make many of our promises - specifically, promises to do things which many others, or we ourselves, have set out and failed to do. But surely *that* can't be right. After all, those promises are some of our most important! A pair of examples will help bring this issue into focus.

Wedding Walt is at his wedding. He is about to promise his partner Wendy to spend the rest of his life with her, for better, for worse. Walt has a lot of evidence that marrying Wendy is the best thing to do: he loves her, and she loves him, and they both very much want to be married to one another. They've been together for a while – through ups and downs – and neither has cold feet. But Walt also has evidence that there's a significant chance that he will not follow through with the promise if he makes it. After all, divorce rates are high, and many apparently equally committed couples have failed to follow through on their vows.³

Presentation Penny and her colleague Pablo are scheduled to make a joint presentation tomorrow morning at work. Pablo asks Penny to promise him that she'll be on time. Penny has a lot of evidence that being on time is the best thing to do: her boss would frown upon her and Pablo for starting their presentation late, Penny gets frazzled when she's running behind and she presents poorly when frazzled, and her tardiness would cause Pablo a lot of undue stress. But Penny also has evidence that there's a significant chance that she will not follow through with the promise if she makes it. After all, Penny is perpetually late, even to big events – like tomorrow's presentation – when it's important for her to be on time.

Suppose the following about these cases: Walt and Penny want to do that which they would be promising to do, and they are volitionally committed to acting as they would promise; they would be making their promises voluntarily and happily; they understand what they would be promising and what it would take

¹ These claims are modified from p. 22 of Marušić (2015).

² The qualification *given that you promise to φ* is a part of the content of the belief necessary for sincerity because, as Liberman (2019) points out (p. 471), “promising to φ usually makes it more likely that you will φ because we are generally motivated to keep our promises.” Importantly, what matters when making a promise is that you believe that you will φ assuming that you promised to φ , even if you wouldn't believe that you will φ assuming, instead, that you didn't promise to φ . That said, for the sake of simplicity, I will sometimes leave this qualification implicit.

³ This example is adapted from Marušić (2015), pp. 3-4.

for them to follow through; following through on the promises is feasible; making and following through on the promises would be good not only for Walt and Penny but also for their promisees; and their promises would be welcomed by their promisees.⁴ In light of this, it seems that Walt and Penny should make their respective promises. In fact, it seems to be not just permissible, but also the best thing for them to do:

But Walt and Penny also have evidence that, nevertheless, there's a significant chance that they won't follow through on their promises if they make them. This evidence comes from a statistical bad track record in Wedding and a personal bad track record in Presentation. The statistical bad track record in Wedding poses a problematic symmetry between Walt and others who set out to spend the rest of their lives with their partners, a significant percentage of whom failed to do so. The personal bad track record in Presentation poses a problematic symmetry between this time and other times when Penny set out to be one time, a significant percentage of which were times when she failed to do so. However significant the chance involved in the relevant problematic symmetry, Walt and Penny have evidence that there's an equally significant chance that they won't follow through on their respective promises.

If this chance is significant *enough*, then Walt and Penny's total evidence wouldn't support the belief that they *will* follow through. Of course, it's hard to pin down exactly how significant a chance needs to be in order to count as significant enough. So, for the sake of our discussion here, let's just pick a number, say 40 percent - if you have an undefeated piece evidence that there is at least a 40 percent chance that you'll fail to φ , then your total body of evidence doesn't support the belief that you will φ .

Suppose, then, that Walt knows that 40 percent of couples who marry end up getting a divorce, and that Penny knows that she runs late around 40 percent of the time when she sets out to be on time. Suppose, too, that Walt doesn't know of anything that distinguishes him from these other spouses-to-be, and that Penny doesn't know of anything that distinguishes this time from these other times. What

⁴ For discussions about the conditions of genuine promising-making, see Fox and Demarco (1996) and Scanlon (1990). See Liberman (2019) for a related discussion on morally permissible promise-making.

should Walt believe about whether he'll stay forever with Wendy, and Penny about whether she'll be on time?

If they believe that they will follow through on their promises, then, according to the evidence claim, their beliefs would be irrational. But if they believe in accordance with their evidence and, thus, don't believe that they will follow through, then, according to the sincerity claim, their promises would be insincere. So – if both the sincerity and evidence claims are true – and if, according to the permissibility claim, a promise should not be made if it would be insincere or if the belief required for sincerity is irrational, then Walt shouldn't promise to spend the rest of his life with Wendy and Penny shouldn't promise to be on time. Call this response – the response that renders the verdict that Walt and Penny (and anybody else) shouldn't promise against the evidence – the hardline response.⁵

The problem is that the hardline response seems to give the wrong verdicts here: it seems that Walt and Penny *should* make their respective promises. This problem becomes more troubling when we step back and see that many other cases fit the same mold as Wedding and Presentation – cases involving promises, perhaps, to quit smoking or to be more attentive to a loved one or to stop procrastinating. These are cases in which it seems that promising is the best thing for the protagonist to do even though the track record of others like them (as in Wedding) or their own personal track record (as in Presentation) suggests that there's a significant chance (40 percent or more) that they'll fail to follow through. Call these *puzzling cases*.

This leave us with a puzzle. We must either reject one of the three claims that generate the problem. Or, in accepting the hardline response, it seems that we must reject the plausible verdicts in puzzling cases – for example, that Walt should promise to spend that rest of his life with Wendy and that Penny should promise to be on time.

⁵ Two terminological notes: (1) The name “hardline response” comes from the idea that those who endorse this response take a hard line on the three plausible claims. (2) The term “promise against the evidence” is shorthand for making a promise when your total body of evidence doesn't support the belief that you will follow through with the promise.

In approaching this puzzle, the most obvious strategy is to solve it by denying one of the three claims that generate it in order to show that it is permissible to promise against the evidence. But this strategy is unsatisfying because, as I will argue below, it is morally impermissible to promise against the evidence. I will then argue for the hardline response by showing that it's not committed to rendering the wrong verdicts in puzzling cases. The hardliner can say that the protagonists in many puzzling cases should make their promises – not because it's permissible to promise against the evidence – but rather because, despite appearances, their total evidence *does* support the belief that they will φ . In many puzzling cases, the argument goes, protagonists' evidence that there's a significant chance that they won't follow through on a promise to φ – the piece of evidence that precludes their total evidence from supporting the belief that they will φ – is defeated, and so, in these cases, the hardliner can say that the protagonists should promise. In puzzling cases where this evidence is not defeated, the hardliner must say that the protagonists should not promise. But, upon closer inspection, this seems to be the right verdict. So, I conclude, the best way to approach the promising puzzle is to show that promising isn't so puzzling after all: the conditions that give rise to the purported puzzle do not often arise and, when they do arise, they are not problematic.

1 Promising Against the Evidence

Denying any one of three original claims amounts to accepting that, sometimes, it's permissible to promise against the evidence. If we deny the evidence claim, then there might be cases where it's rational for you to believe that you will φ even though your evidence doesn't support that belief. So, in the puzzling cases, it would be permissible for the protagonists to make their promises even though their beliefs that they will φ aren't supported by the evidence, as long as their beliefs meet whatever rationality conditions are applicable.

If we deny the sincerity claim, then there might be cases where your promise is sincere even though you don't believe that you will φ . So, in puzzling cases, it would be permissible for the protagonists to promise even when they don't believe that they will φ , as long as they have whatever attitude is necessary for sincerity.

And, finally, if we deny the permissibility claim, then there might be cases where you should make a promise even if either the promise wouldn't be sincere or the belief required for sincerity isn't rational. So, in puzzling cases, it would be permissible for the protagonists to make their promises even if they don't believe that they will φ , or even if they do believe that they will φ but their belief isn't rational.

Berislav Marušić develops a solution to the promising puzzle which rejects the evidence claim,⁶ and Alida Liberman, as I interpret her, would reject the permissibility claim in favor of subtly different alternative.⁷ Their respective approaches are worth considering at length.

Marušić's solution starts off with a distinction between two ways that you can settle the question of whether you will φ : a theoretical way and a practical way. In order to settle the question in the theoretical way, you must look to your evidence and engage in theoretical reasoning – you must weigh your evidential reasons bearing on whether it's true that you will φ . If your evidence suggests that you will φ , then you should believe that you will φ . If it doesn't, you shouldn't. The outcome of theoretical reasoning is a prediction about whether you will φ .

In order to settle the question of whether you will φ in the practical way, you must first answer the question of whether you ought to φ by engaging in practical reasoning: you must weigh your practical reasons for and against φ ing. The outcome of practical reasoning is a decision whether to φ . If your reasons tip in favor of φ ing, you should decide to φ and if they don't, you shouldn't. By deciding whether to φ , you indirectly settle the question of whether you will φ . If you decide to φ , the answer is yes, you will φ ; if you decide not to φ , the answer is no, you won't.

When it comes to promising, Marušić goes on to argue, you should settle the question of whether you will φ the practical way – and not the theoretical way. So, according to Marušić, Walt shouldn't look at his evidence and predict whether he'll spend the rest of his life with Wendy. Instead, he should weigh

⁶ Marušić (2013, 2015)

⁷ I add "as I interpret her" here, not because of any lack of clarity in Liberman, but because she does not directly address the promising puzzle; instead, she argues for several necessary conditions for permissible promise-making and then considers whether promises made under uncertainty can satisfy those conditions. Because of this, my take on Liberman's approach to our puzzle involves extrapolation.

his practical reasons for and against doing so and thereby settle, and thus come to believe, that he will. Similarly, Penny shouldn't look at her evidence and predict whether she'll be on time. Instead, she should weigh her practical reasons for and against doing so and thereby settle, and thus come to believe, that she will.

Since you settle the question in a practical way in cases of promises, Marušić argues, your belief about whether you will φ is subject to the standards of practical rationality only; theoretical standards don't apply. In puzzling cases, the protagonists' beliefs that they will φ are rational in light of the applicable practical standards of rationality even though they would be irrational in light of the (inapplicable) theoretical standards. In these cases, Marušić concludes, it's permissible for the protagonists to promise against the evidence – after all, they have the belief necessary for promissory sincerity and that belief is rational.

Now let's turn to Liberman. She starts by developing four conditions of permissible promise making, one of which regards the reasonableness of promisers' beliefs that they will follow through on their promises. This condition – she calls it *the realistic self-assessment condition* – is in the same vein as the condition articulated in the permissibility claim: “A promise is permissible to make only if the promisor *blamelessly* believes that she will be able to carry out the promised action.”⁸

Liberman goes on to specify that whether you blamelessly (rationally) believe that you will follow through on your promise depends on whether this belief is reasonable in light of a realistic assessment of “yourself and your abilities.” This *self-assessment*, she then clarifies, takes into account the various internal factors that determine how resolved you are follow through, such as the strength of your will and your susceptibility to temptation.⁹ But it does not take into account the various external factors or “outside sources” beyond your control that may prevent you from φ ing but that would excuse you from the moral censure that normally accompanies failing to fulfill a promise – for example, a car accident that prevents you from fulfilling your promise to pick up your friend at the airport. It will be helpful

⁸ Liberman (2019), p. 473

⁹ Liberman (2019), p. 481

here to see Liberman as endorsing a qualified version of the permissibility claim: it is permissible for you to promise to φ only if your promise is sincere and the belief required for sincerity is rational *given the subset of your evidence regarding the relevant internal factors*.

Endorsing this qualified version of the permissibility claim allows Liberman to say that it is permissible to promise in puzzling cases in which the reason why there's a 40 chance you won't follow through has to do to external, rather than internal, factors. This entails that, in some cases, its permissible to promise against the evidence – namely, when your belief that you will φ is supported by the subset of your evidence regarding the relevant internal factors, but not supported by your total evidence, including considerations about external factors.

2 A Worry for Promising Against the Evidence

Each solution that denies one of the three claims faces its own individual worries.¹⁰ Instead of discussing them all in turn, let's focus on one general worry that troubles any such solution: it underrates the moral importance of a promiser's evidential situation.

To feel the force of this worry, it's helpful to think about the nature of promissory obligation. A promissory obligation is the moral obligation to φ that you incur when you promise to φ .¹¹ Consider the expectational account.¹² It says that you're morally obligated to φ when you promise to φ because promising to φ leads your promisee to form an expectation that you will φ . Breaking a promise is tantamount to deceiving your promisee and, because deceiving is wrong, breaking promises is wrong, too. Another account says that you're morally obligated to φ when you promise to φ because, by promising to φ , you invite the promisee to trust you to φ .¹³ It is wrong to break promises because it amounts to a wrongful breach of trust.

¹⁰ For worries and criticisms of Marušić's view, specifically, see Vavova (2018).

¹¹ For helpful surveys and discussions of accounts promissory obligations, see Downie (1995), Shiffrin (2008), Friedrich and Southward (2011), Owens (2012), and Liberto (2016).

¹² Scanlon (1990) and Scanlon (1998), chapter 7

¹³ See Friedrich and Southward (2011) for the trust-based account. Judith Jarvis Thomson develops a similar reliance-based account according to which that you're morally obligated to φ when you promise to φ because, by promising to φ , you invite the promisee to rely on you to φ . See pp. 294 – 321 of Thomson (1990).

If something along these lines is right, there seems to be a serious problem with promising against the evidence. The problem is that it's wrong to foster in someone the sort of expectation or trust characteristic of promissory obligation if your evidence doesn't suggest that you will do what that person is expecting or trusting you to do. If you promise to φ when your evidence suggests there's a significant chance you won't φ , you put yourself at a correspondingly significant risk of deceiving your promisee or breaching their trust. If you know there's a significant chance of committing a moral wrong by doing something – in this case, by promising against the evidence – then that's surely a weighty moral consideration against doing that thing.

There's another related, perhaps even more serious, problem with promising against the evidence. When you promise to φ , you give the person you promise evidence that you will φ . Not only that, but you also invite them to take into account that you will φ when they make decisions and plans. After all, a promise to φ involves more than just the corresponding assertion. If Walt promises Wendy that he'll spend the rest of his life with her, for example, he's in effect saying, "Rest assured, Wendy, I will spend the rest of my life with you. You can count on it." It is in light of this that Wendy will expect (trust) him to do so and it is this expectation (trust) which will then inform her decisions and plans.

But what decisions and plans are best for your promisee to make can depend heavily on whether you will φ . If Walt will spend the rest of his life with Wendy, for example, it'll be best for Wendy to make one set of plans and decisions: to pass up a promotion at work to spend more time with Walt, to have a child with him, to move to the suburbs. If Walt won't, however, it may be best for Wendy to make a different set: to take the promotion, to remain childless, to stay in the city she loves. Generally, if your promisee's evidence suggests that you will φ , it may be reasonable for them to make the set of plans and decisions that are best given that you will φ ; but if your promisee's evidence doesn't suggest this, it may be reasonable for them to make a different set. When you promise against the evidence to φ , you give your promisee evidence that makes it the case that it's reasonable for them to make the set of decisions and plans that are best given that you will φ , even though your own evidence may suggest that it's reasonable for them to make those in the other set.

Promising against the evidence is especially problematic in cases like Wedding, when the resulting expectation (trust) affects important, life-altering decisions and plans of your promisee. There's something reckless and selfish about Walt promising Wendy that he'll spend the rest of his life with her – knowing that Wendy will make important decisions and plans, and perhaps great personal sacrifice, in light of his promise – when his own evidence doesn't suggest that he will.¹⁴ But the problem remains in cases when there's less at stake. At the very least, you disrespect your promisee's time and interests when invite them to make decisions and plans that your own evidence suggests may not be the best for them to make.¹⁵

This worry for promising against the evidence gets a couple of issues at the heart of morally permissible promise-making. First of all, what matters for promise-making is not, primarily, that the belief necessary for sincerity is *rational*, but rather that it is supported by the evidence. The importance of the belief's being rational is contingent on the rationality of belief being the same as the belief being supported by the evidence – that is, on the truth of the evidence claim. Solutions like Marušić's go wrong because they assume that rationality – however it is cashed out – is primarily important. But, as this worry brings out, if the protagonist's belief that he will φ is “rational” but not supported by the evidence, then it is not morally permissible for him to promise.

¹⁴ Considering the various ways that societal structures of work and family disadvantage women – and the disproportionate harms that mothers tend to incur in divorce as a result – sheds even more light on the moral importance of promises, particularly on the promises of husbands-to-be in cases like Wedding that involve heterosexual couples intending to have children. As Saul (2003) explains, society is set up so that the best jobs – secure, high-paying full-time jobs – make demands that are extremely difficult for primary childcare-givers to satisfy. Mothers tend to take on most of a family's childcare responsibilities, which allows fathers to take advantage of the opportunities of the best jobs but leaves mothers financially dependent on their husbands, since most jobs whose demands they *can* meet are low-paying and part-time. So it's no surprise that divorce tends to leave mothers in dire financial straits. Divorced mothers remain unable to meet the demands of the best jobs since they typically have primary custody. What's more, they no longer have access to their partner's income since courts rarely grant alimony and child support is often meager; they also now must bear the burdens, financial and otherwise, of single parenthood. Saul notes that, because of all this, 45 percent of households of divorced mothers are at or below the poverty line. See chapter 1 of Saul (2003), especially pp. 8-11. See Okin (1989) for further discussion.

¹⁵ Liberman (2019) makes a similar point on p. 471: “it is morally impermissible to allow or actively lead people to form false beliefs, especially about matters that are important to them or relevant to their interests.” It's worth noting that she goes on to consider ways in which some promises made under uncertainty – including some promises against the evidence – can avoid being morally impermissible in this way. For the relevant discussion, see Liberman (2019), p. 480-481. Ultimately, I do not think that these considerations absolve the relevant promises against the evidence, but will not go into why here – doing so would require a lengthy discussion on subtly different assumptions Liberman and I make about deception, promises, and promissory obligation that would take us too far afield.

Secondly, this worry clarifies that, when it comes to permissible promise-making, what matters is that your belief you will follow through is supported by *all* of your evidence, not just a particular subset. After all, the relevant risks and harms that render promising against the evidence morally impermissible obtain *whenever* your evidence suggests that there's a 40 percent chance that you will not follow through, no matter if it's an internal lack of resolve or external circumstances that makes it the case that there is such a chance.¹⁶

3 What We Want Out of an Approach

If an approach to the promising puzzle is to avoid this worry, it must be some type of hardline response – that is, it must deny that it is sometimes permissible to promise against the evidence. Of course, hardline responses also face their own worry: they seem to render the wrong verdicts in puzzling cases. A good hardline response, then, must address this worry.

There are two broad ways for a hardline response can do this. It can account for why, despite our initial verdict, protagonists in puzzling cases really shouldn't promise; or, it can show how the hardliner can consistently say that protagonists in these cases should promise. The second way is better than the first. The promises involved in puzzling cases are some of our most important, woven deeply into our relationships and social institutions – a good hardline response can't issue a blanket prohibition on making them.¹⁷ This means that a good hardline response must show that, in many puzzling cases, the protagonists' belief that they will φ is supported by their total body of evidence, despite them having a piece of evidence from either a statistical or personal track record that there's a significant chance that they won't.

¹⁶ On Liberman's view, that there are a number of reasons why it matters whether it's internal or external factors that make it the case that there's a significant chance that a promiser will not follow through. Specifically, Liberman argues that it is difficult for promises made under internal uncertainty, but not promises made under external certainty, to fulfill the good faith condition – another of her four necessary conditions for permissible promise-making. See section 9 of Liberman (2019). I think she is probably right about this. So, to clarify, I am not denying that whether a promise is made under so-called internal or external uncertainty can make a difference to its permissibility; I am just denying that it explains why it can be morally permissible to promise against the evidence.

¹⁷ I assume that, in order to count as a good response, a hardline response needn't show that the protagonists in *all* puzzling cases should make their promises – just in many. But the hardliner should be able to draw a principled distinction between puzzling cases in which it is permissible for the protagonists to promise and the ones in which it isn't.

What else do we want out of an approach to the promising puzzle? In the process of proposing his solution, Marušić levels several criticisms against hardline responses. Since we know we want a hardline response, responding to these criticisms forms the basis of a couple of additional desiderata.

Marušić's first criticism of hardline responses is that they promote what he calls epistemic evasion. Epistemic evasion happens when you refuse to commit to φ ing – even when φ ing is clearly the best thing for you to do – just because your evidential considerations don't allow you to predict that you will φ . In Marušić's words:

We use epistemic evasion to refuse to take responsibility for our actions. We refuse to commit to doing something, and we take the mere fact that there is evidence that we would fail to follow through if we decided to do it as an excuse—like someone who refuses to commit to being faithful just because he has evidence that there is a significant chance that he will fail. Then our failure is moral and we are blameworthy. Someone who claims to just be an unfaithful person—pointing contritely, but with quiet pride, to a track record of broken trust—is indeed a bastard!¹⁸

Another thing we might want out of a hardline response, then, is that it avoid recommending epistemic evasion, or else it explain why it isn't morally problematic.

Marušić also criticizes hardline responses because he claims that they often can't account for the asymmetry between what a promiser should believe and what a detached observer should believe about whether the promiser will φ .¹⁹ To motivate the thought that there should be such an asymmetry, Marušić starts off by arguing that whether you should believe that you will φ depends in part on whether φ ing is up to you. He compares two similar cases of belief. Suppose that you want to complete a 5k and a marathon this summer. The 5k is popular, and there's a lottery that determines who gets to run it – 40 percent of the people who enter the lottery don't win. The marathon, on the other hand, admits anybody who wants to run it but it is very difficult - 40 percent of people who are admitted fail to complete it.

Before you know the results of the lottery, it seems that you shouldn't believe that you will complete the 5k: this is a significant chance that you won't get in and also whether you do is not up to you. But it seems that it would be permissible for you to believe that you will complete the marathon: there is a significant chance that you won't complete it but whether you do *is* up to you. That said, Marušić argues,

¹⁸ Marušić (2015), pp. 119-120

¹⁹ Marušić (2015), pp. 20-22

a detached observer should have the same belief in both cases about whether you will complete the race in question since the chance that you won't complete it in each case is equally significant.

In all cases of promising, the relevant belief is about whether the promiser will do something that's up to them. So promisers' beliefs about whether they will φ should be different from a detached observer's belief about what they will φ – even if the promiser and detached observer have the same evidence. Walt's belief about whether he will spend the rest of his life with Wendy should be different from a sociologist's belief about the same thing, for example, even though both Walt and the sociologist have access to all the same information about divorce rates.

Marušić argues that hardline responses often can't account for this asymmetry, including in puzzling cases – after all, both the protagonist and a detached observer can have access to the same bad track record which would give them each evidence precluding their respective total bodies of evidence from supporting the belief that the protagonist will φ . According to the claim about evidence which the hardliner accepts, then, both the protagonist and a detached observer shouldn't believe that the protagonist will follow through. So something else that we might want out of a hardline response is that it can account for the asymmetry which Marušić says it can't.

4 The Hardline Response

Our assumption from the outset has been that if you have a piece of evidence that there's a 40 percent chance that you won't φ , then your total body of evidence doesn't support the belief that you will φ . But, of course, the assumption holds only if this piece of evidence is undefeated. The hardline response I want to develop centers on the claim that, for protagonists in many puzzling cases, this piece of evidence *is* defeated.

This hardline response starts off with a natural thought: sometimes, your evidence supports the belief that you will φ even though many others, or you yourself, have set out and failed to φ . It's easy to come up with examples: Serena Williams' evidence supports the belief that she will qualify for Wimbledon even though many others before her tried and didn't qualify; the evidence of a high-polling candidate for congress may support the belief that she will win the election even though many others

ran for that seat and lost; a student's evidence may support the belief that he will pass his math exam on Monday – he's started seeing a tutor – even though he has failed math exams many times before.

In each of these cases, the person in question has good reason to believe that they are not relevantly like the other people, or their past self at other times, who set out to φ but failed, and that's why their chance of success cannot be read off the success rate of those other people or times.

The hardline response says that this is what happens in many puzzling cases: the protagonists have good reason to think that they are not, in fact, relevantly like the comparison group - the other people cited by statistical bad track records or their past self at the other times cited by personal bad track records. This good reason defeats the piece of evidence they get from the bad track record which cites that comparison group. After all, if the protagonists have good reason to believe that they are not relevantly like the comparison group, then it is not reasonable for them to infer that their chance of success mirrors the success rate of the comparison group and, thus, the track record no longer provides them with grounds for believing that there's a significant chance that they won't φ .

With that piece of evidence defeated, the protagonists' total evidence is no longer precluded from supporting the belief that they will φ . If the protagonists' total evidence now *does* support that belief, and if the protagonists so believes, they have the belief that is necessary for promissory sincerity and that belief is rational. In these cases, the protagonists' promises satisfy the conditions of the permissibility claim and so the hardliner is in a position to render the right verdicts – the protagonists *should* make their promises. My task in the remainder of the paper is to look more closely at puzzling cases of this sort and, in the process, explain why we should think that many puzzling cases *are* of this sort and why we should accept the hardliner's verdicts in the puzzling cases that aren't.

5 The Hardline Response: No-symmetry Cases

Let's first look at puzzling cases in which the protagonists have a good reason from the start to believe that they are not relevantly like the comparison group. Call these no-symmetry cases: the protagonists' evidence suggests that the problematic symmetry posed by the relevant track record doesn't hold to begin with. In these cases, the protagonists have evidence about themselves or their circumstances that

suggests that that they are not relevantly like the comparison group. Broadly, there are two main ways this evidential situation can play out.

The first way is that the protagonists have evidence suggesting that they are much more likely than the comparison group to φ . This may be because the protagonists can point to a specific feature that distinguishes them in this way from the comparison group. Suppose that Penny knows that she cares a great deal about being on time tomorrow: in addition to being her colleague, Pablo is a dear friend, and Penny is eager to avoid stressing him out. Or suppose, for another example, that a protagonist has tried and failed to quit smoking many times before but, this time, a baby is on the way.

Or this may be because the protagonists have evidence about themselves or their circumstances which suggests they are very likely to φ , and, from this, they reasonably infer that they are much more likely to φ than the comparison group. Suppose that Walt knows that he and Wendy have the same core values, compatible personalities, and complementary strengths and weakness; he desires being in a healthy marriage over pretty much everything else; he and Wendy have built a stable, loving relationship over the years; both sets of their parents have fostered flourishing marriages; and – to top it off – they belong to a strict religious community whose members rarely ever cheat or divorce. Walt knows that all of this makes it very likely that he will spend the rest of his life with Wendy, and, thus, he reasonably infers that he is much more likely to do so than members of the comparison group, whose average success rate of staying together is mediocre. Note that this inference is reasonable even if Walt does not know much at all about the members of the comparison group besides their average success rate.

The second way that this evidential situation can play out is that the protagonists have evidence suggesting that the features they share with the comparison group are not telling of their chance of success at φ ing. In order to be *relevantly* alike, the protagonist and comparison group must share *relevant* features – namely, features that are the same as, or related to, the features that are causally responsible for their respective chances of success. They are not relevantly alike if the shared features are something irrelevant, like favorite ice cream flavor or hair color. After all, if what's causally responsible for the comparison group's success rate is not the same as, or related to, what the group and protagonist have

in common, then it would be a mere happy accident if the group had an average success rate that reflects the protagonist's chance of success.

Suppose that Walt has a decent grasp on the features – the values, interests, desires, goals, personality traits, circumstances, and so on – that go into determining how likely it is that a marriage will last. Suppose, too, that he has evidence – gathered from a lifetime of interacting with a plethora of people in his demographic group – which suggests that couples demographically similar to him and Wendy vary greatly when it comes to those features. Since demographic features are not causally responsible for making a marriage last, nor are they correlated with the features that are, Walt has good reason to believe that he is not relevantly like the comparison group.

To sum up so far: in no-symmetry cases, protagonists have good reason from the outset to think that they are not relevantly like the comparison group. This may be because they have evidence that they are much more likely than the comparison group to φ , or because they have evidence that the features they share with the comparison group are not very telling about their chance of success. In all these cases, the evidence that precludes the protagonists' total evidence from supporting the belief that they will φ is defeated. Assuming that the protagonists' total evidence now does support this belief, and that they so believe, the hardliner is in a position to render the right verdict: they should make their promises.

In no-symmetry cases where the total evidence does not now support this belief, the hardliner must say that protagonists shouldn't promise. But this isn't a ding against the hardliner. Suppose, for example, that Penny can point to a feature making her much *less* likely than normal to be on time. Although this is sufficient for a good reason to believe that she is not relevantly like the comparison group, Penny's evidence does not now suggest that she will be on time. So, in this case, the hardliner must say that Penny should not make her promise. And this seems to be the right verdict.

6 The Hardline Response: Broken Symmetry Cases

Now let's look at puzzling cases in which the protagonists do not start off with a good reason to believe that they are not relevantly like the comparison group. Even in these cases, their total body of evidence can end up supporting the belief that they will φ if they gain such a reason in the process of deliberating

about whether to promise. Call these broken symmetry cases: the protagonists' evidence initially suggests the problematic symmetry holds but they get a new evidence suggesting it doesn't.

In some broken symmetry cases, protagonists gain the sort of evidence that they had from the start in no-symmetry cases. For example, suppose that Walt introspects and identifies the desires, values, beliefs, and goals that make him very likely to spend the rest of his life with Wendy only after being confronted with the discouraging statistics about divorce.

But, I want to argue that, in other broken symmetry cases, this evidence comes from a special source – namely, a special sort of seeming or inclination. Imagine a protagonist in a statistical bad track record puzzling case who is deliberating about whether to promise to do something that many others like him have set out and failed to do. When confronted with the track record, it strikes him that he is more committed or has more determination or wants it more than they do – *he is different* from them. Or imagine a protagonist in a personal bad track record puzzling case who is deliberating about whether to promise to do something that she, herself, has set out and failed to do many times over. When confronted with the track record, it strikes her that she is more committed or has more determination or wants it more this time – *this time is different*. These are the special sort of seemings in question. Call them *difference-seemings*.

A few general points about seemings are in order.²⁰ Roughly, a seeming is a phenomenologically-laden impression that the proposition you are considering is true. The phenomenology characteristic of seemings is striking-as-true. Seemings and beliefs often go together. When you believe, say, that you will complete a marathon, it might also strike you as true that you will complete a marathon. But seemings and beliefs can come apart. You can have a seeming that something is true without the corresponding belief and vice versa. This is important because, in broken-symmetry cases, protagonists have a difference-seeming without the corresponding belief. In fact, the difference-seeming ends up providing the grounds for the corresponding belief.

²⁰ If you don't like talk of seemings, you can instead think of difference-seemings as difference-inclinations or difference-hunches: the protagonists have an inclination / hunch that it's true that they're different or that this time is different.

Seemings come in degrees of strength. In other words, a proposition can strike you more or less as true. Your seeming that $2+2=4$, for example, is probably stronger than your seeming (supposing you have one) that your favorite normative ethical theory is correct. This is important because, for a difference-seeming to be worth much, it must be fairly strong. It also must be fairly steady. A difference-seeming won't count for much if it flickers in and out.

With that, we get to the main contention here: the presence of a strong and steady difference-seeming can sway the permissibility status of a promise in puzzling cases. Consider the following variations on Wedding and Presentation.

Wedding Variations Walt knows that spending the rest of his life with Wendy – and promising her to do so – is the best thing for him to do. But he also knows that 40 of couples demographically similar to him and Wendy wind up divorced. Walt doesn't have any initial reason to believe that he is relevantly different from demographically similar spouses-to-be, many of whom failed to follow through on their vows. **Variation 1** Nevertheless, it strongly *seems* to Walt that he is more committed – whenever Walt contemplates his situation, it strikes him that he is special, that his level of commitment to Wendy far outruns that of the average millennial. **Variation 2** Moreover, it doesn't seem to Walt that he's any more committed – whenever he contemplates his situation, either he doesn't have a seeming regarding whether he is more committed, or he does have a seeming, but his seeming is that he's not any more committed.

Presentation Variations Penny knows that being on time to the presentation – and promising Pablo to be on time – is the best thing to do. But she also knows that she runs late 40 percent of the time when she sets out to be on time. Penny doesn't have any initial reason to believe that there's something about her or her circumstances that makes her more likely than normal to be on time. **Variation 1** Nevertheless, it strongly *seems* to Penny that she is very determined to be on time tomorrow – when Penny contemplates her situation, it strikes her that this time is special, that her level of determination far outpaces her normal level. **Variation 2** Moreover, it doesn't seem to Penny that this time's any different – when she contemplates her situation, either she doesn't have a seeming regarding whether she's very determined, or she does have a seeming, but the seeming is that she's not any more determined than normal.

The main contention, as applied to these cases, is that the protagonists should make their promises in the first variations, but not in the second variations, and the reason for this difference in verdicts is the presence, or lack thereof, of a difference-seeming. Difference-seemings make a difference to the rationality of the protagonists' belief that they will φ : the belief is rational, and therefore the promise is permissible, in the first variations but not in the second variations.

Of course, this contention is plausible only if difference-seemings are epistemically significant. The idea is that when the protagonists have a strong and steady difference-seeming, they gain *evidence* that

they are, in fact, different from the other people cited by a bad statistical track record or that this time is, in fact, different from the other times cited by a bad personal track record. This evidence gives the protagonists good reason to believe that they are not relevantly like the comparison group, thus breaking the problematic symmetry posed by the bad track record.

The idea that seemings are epistemically significant isn't new. Phenomenal conservatives have been defending that idea for a long time – and apply epistemic significance to seemings across the board.²¹ But phenomenal conservatism is not uncontroversial. It would be better, then, to have an account of the rational basis of difference-seemings that's acceptable to those who reject phenomenal conservatism. Such an account would just need to explain why difference-seemings, in particular, have a rational basis.

Here's a suggestion: difference-seemings generally arise from internal features that determine the truth of their propositional content. So, for example, Walt's and Penny's difference-seemings arise from affective and motivational states – the various cares, values, desires, goals, and so on – that go into determining their high levels of commitment and determination, respectively.

Taking this suggestion on board helps clarify the role difference-seemings play in broken symmetry cases. For whatever reason, protagonists in these cases cannot put their finger on the features that give rise to their difference-seemings. Perhaps they simply haven't introspected, or perhaps they have introspected but these features are occluded from them – they just can't get a good read on them or they downright misread them. Their difference-seemings, then, serve as a proxy for knowledge of these features. They indicate to the protagonists that there is something about them that makes them, for example, very committed or very determined, even though they can't place their finger on what, exactly, it is.

But, notably, the seeming isn't just they are very committed and very determined – it's that they are *more* committed and *more* determined. Difference-seemings, then, catch an implicit inference from high

²¹ Phenomenal conservatives defend principles like this: if it seems to you that p, then, in the absence of defeaters, you thereby have at least some positive justified reason for believing that p. For defenses of this principle (or principles close to it) see Huemer (2001, 2006, 2013), Lycan (1988, 2013), Skene (2011), and Tucker (2010).

levels to *higher* levels – Walt’s inference that his high level of commitment must be higher than the other spouses-to-be cited who have a mediocre success rate of staying together and Penny’s inference that her high level of determination this time must be higher than at the other times cited given her mediocre success rate of being on time.

But why think that this suggestion is correct? Whether difference-seemings do, in fact, generally arise from the relevant truth-related features is an empirical question. Psychological studies would need to test whether promisers who report having a strong and steady difference-seeming before promising are more likely than those who report otherwise to follow through on their promises. If these studies showed that the rate of follow-through was significantly higher among promisers with difference-seemings than without, we would have empirical proof that difference-seemings are generally grounded in that which makes you likely to follow through.

That said, if we assume some sort of internalism about epistemic rationality, it doesn’t have to be the case that difference-seemings do, in fact, generally arise from truth-related features in order for them to be epistemically significant; it just has to be the case that the suggestion that they do is plausible. This is important because it means that we don’t have to wait on empirical work to know whether difference-seemings are epistemically significant. We just have to look and see whether the suggestion is plausible given the evidence currently available. So, is it plausible? I think it is.

It’s helpful here to reflect on some familiar personal experiences. Consider times when it seems to you that you are more or less likely than normal to do something even though you don’t know why. It may seem evident to you, for example, that you will not get on the Ferris Wheel at the fair with the rest of your friends even though it’s never dawned on you that you’re acrophobic. Or it may strike you as true that you will join in on weekly drinks with your colleagues, which you normally pass on – even though you don’t realize it, you are avoiding the mountain of grading waiting for you at home. Or consider times when you have a seeming about what you would do in some hypothetical dangerous scenario. It might seem obvious to you would run away in a frenzy (or that you would confront the danger head-on) even though you aren’t aware that you’re prone to panic (or admirably brave). I take it

that such experiences are familiar, which implies that seemings about prospective actions can and do arise from unknown features that help determine the truth of their propositional content. And, if that's the case, then the suggested explanation of why difference-seemings have a rational basis, far from being ad hoc or new-fangled, seems plausible.

So, we do have some reason in favor of the suggestion. One might worry, though, that existing empirical work on self-serving cognitive biases gives us better reason to think that difference-seemings arise from such biases. The phenomenon of illusory superiority, for example, is well-documented in psychology.²² Studies show that people typically report that they have positive characteristics and skills to a higher degree, and negative characteristics and ineptitudes to a lower degree, than the average person. The thought, then, is that it's more likely that difference-seemings arise, not from relevant truth-related features, but rather from self-serving cognitive biases like illusory superiority.

Thankfully, this worry is less troublesome than it may first appear, for a couple of reasons. First, illusory superiority has to do with comparison to others, so it could only give rise to difference-seemings in puzzling cases involving statistical bad track records. Second, the relationship between the relevant empirical work and difference-seemings is too loose to safely infer from this work that even *these* difference-seemings likely arise from illusory superiority. Puzzling cases involve a seeming that's triggered when protagonists contemplate a particular action along with how likely they are to do that action at some point in the future. The relevant empirical works shows that people tend to rate themselves better than the average person when faced with a list of amorphous character traits or skills, untethered to any particular action. It would be a far stretch infer from these studies that – in the specific circumstances of puzzling cases – illusory superiority likely causes it to seem to protagonists that they are more committed or determined to φ , for example, than they actually are.

This stretch becomes even further when considering the relevant sense in which protagonists who are more determined or committed to φ are superior to the comparison group: they are superior in that they are more likely to φ . But this narrow sort of superiority doesn't seem to be closely related to

²² Alicke (1985) and Hoorens (1993)

superiority of character or skill that illusory superiority tends to distort when you evaluate yourself.²³ It's doubtful, then, that a tendency to see yourself as superior without warrant when it comes to amorphous character traits and skills would cause you to see yourself as superior without warrant in this narrower way related to your likelihood to φ .²⁴

Where, then, should we come down on the epistemic significance of difference-seemings? If pheromonal conservatism is true, then difference-seemings – like all seemings – are epistemically significant. But, if pheromonal conservatism turns out to be incorrect, difference-seemings are still epistemically significant just in case the suggestion that they arise from truth-relevant features is plausible. Even though we lack empirical proof for this suggestion, reflecting on familiar personal experiences lends does lend this suggestion some plausibility, and existing psychological studies on self-serving cognitive biases do not undermine it. It seems, then, that there's some – though defeasible and not decisive – reason to accept difference-seemings as epistemically significant.²⁵

In sum: in broken symmetry cases, the protagonists do not start off with good reason to believe that they are not relevantly like the comparison group, but they gain it in the process of deliberating whether to promise. In some broken symmetry cases, it's been argued, the evidence that gives them this good reason is provided by a difference-seeming. Assuming that difference-seemings are epistemically significant, they give protagonists evidence that they are more likely to φ than the other people cited by a statistical bad track record or their past self at the other times cited by personal bad track records. This evidence provided by difference-seemings defeats the evidence provided by the track record that there's

²³ It's true that having certain character virtues to a high degree – having a lot of grit, for example, or being very trustworthy – may make it more likely for you to follow through on any given promise. But, generally, your level of commitment or determination to φ depends heavily on value-neutral features of your personality and circumstances – your particular preferences, priorities, and goals. This is why it's doubtful that a (perceived) high level of determination or commitment to φ is reliably correlated with (perceived) superiority in character or skill.

²⁴ I've focused on the limits of empirical studies on illusory superiority since it seems to be especially pertinent to this discussion. That said, I assume that similar things can be said about the limits of empirical studies on other self-serving cognitive biases.

²⁵ In section 8, there is a discussion about what it means for the hardline response if we learn that difference-seemings do not, in fact, arise in the suggested way, in which case there would no longer be reason to accept them as epistemically significant.

a significant chance they won't φ , thus allowing the protagonists' total evidence to support the belief that they will φ .

7 What About Untrustworthy Seemings?

It might be worried that protagonists' difference-seemings can arise from wishful thinking or self-deception, in which case it would not rationalize their belief that they will φ . We might think of friends who consistently promise and fail to return borrowed books or pay us back for dinner, and also of more serious cases – an absent father, for example, who consistently fails to follow through on promises to show up to his child's parties and plays and recitals.

It seems that the hardliner is committed to saying that, if these people have a difference-seeming, their beliefs that they will φ is rational. This would be a troubling commitment, especially if we know that the absent father, for example, has a history of untrustworthy difference-seemings when he contemplates whether he will show up – surely it's not rational for him to believe that he will this time.

But the hardliner isn't so committed. That's because the evidence that difference-seemings provide can, itself, be defeated. Suppose, for example, that Penny regularly has difference-seemings when she contemplates whether she'll be on time yet still fails to be on time. Even when it seemed to her in the past that this time is different, she continually proceeded to run late. In this case, Penny has evidence that her difference-seeming does not indicate that she is, in fact, more likely than normal to be on time. The evidence that gives her reason to believe that she is not relevantly like the comparison group provided by such a difference-seeming, then, would be defeated.

Of course, the evidence provided by difference-seemings can be defeated by things other than a history of untrustworthy ones. Suppose that, in addition to his strong and steady difference-seeming, Walt knows something about himself that makes him very unlikely to spend the rest of his life Wendy – say he knows that he has a strong distaste for monogamy that Wendy doesn't share. In this case, his evidence of his attitude towards monogamy defeats the evidence from the difference-seeming that he is more committed than the average spouse-to-be.

In broken symmetry cases where the evidence provided by a difference-seeming is defeated, it is not rational for protagonists to believe that they will φ - even if they have the seeming in full force. If the evidence provided by the difference-seeming is defeated, it doesn't serve its own defeating purpose - it doesn't defeat the evidence from the relevant bad track record that there's a significant chance that they won't φ . So, in these cases, the protagonists' total evidence remains unable to support the belief that they will φ and, thus, the hardliner must say that they shouldn't make their promises.

But this seems to be just the right verdict. It seems that Penny should not promise, and neither should Walt. More generally, it isn't permissible for any protagonist with a history of untrustworthy difference-seemings - the book-borrower, the free-riding dinner companion, or the absent father - to make promises based on another difference-seeming.

8 Evaluating the Hardline Response

Let's conclude by looking at whether the hardline response gives us what we want out of an approach to the promising puzzle. One thing we want is that it avoid the general worry troubling solutions that entail that it is sometimes permissible to promise against the evidence. The worry was that, by promising against the evidence, you take on a high risk of deceiving, or breaching the trust of, your promisee; relatedly, you invite your promisee to make decisions and plans based on the assumption that you will follow through, even though your own evidence doesn't suggest that you will. The hardline response avoids this worry because it maintains without concession the impermissibility of promising against the evidence: if your total body evidence does not support the belief that you will φ , you should not promise to φ .

Another thing we want out of a hardline response is that it render the right verdicts in puzzling cases. Puzzling cases are those in which it seems that the protagonists should promise against the evidence. The hardline response shows that, despite first appearances, in many puzzling cases, the protagonists' total evidence *does* support the belief that they will φ , and so the hardliner can say that they should promise. The hardliner renders the right verdicts in these cases.

Now, it's true that there are some puzzling cases where the hardliner cannot say that these protagonists should promise. These include puzzling cases like the second variations of Wedding and Presentation – cases in which protagonists do not have good reason from the start to believe that they are not relevantly like the comparison group, nor do they gain it. But the hardliner gets it right here: it would be impermissible for Walt and Penny to promise in the second variations of their respective cases.

The hardliner must say that protagonists shouldn't promise in a few more types of puzzling cases, too – namely, in no-symmetry cases where protagonists have good reason to think that they are not relevantly like the comparison group but in ways that make them less likely to φ , and also in broken symmetry cases where the evidence provided by the difference-seeming is defeated. But, as we've seen, this seems to be the right thing to say. So, both in puzzling cases where protagonists should promise and also in puzzling cases where, upon closer inspection, it's clear they shouldn't, the hardline response renders the right verdicts.

It's also worth addressing what happens to the hardline response if psychology shows us one day that difference-seemings do not generally arise from features that help determine the truth of their content, in which case the suggestion that they do loses its plausibility and difference-seemings would no longer be epistemically significant. The downside would be that, on the hardline response, there would be considerably less cases in which promising is permissible. Perhaps this would make the hardline response a bit less satisfying, but it's worth emphasizing that the hardliner could still say that the protagonist should promise in a large number of puzzling cases – namely, no-symmetry cases and broken-symmetry cases where the relevant evidence is provided by something other than a difference-seeming.

More importantly, the hardliner seems give the right verdict here, however unsatisfying. In learning that difference-seemings do not arise from truth-related features, we would learn something morally important, albeit unhappy, about ourselves that we should take into account when making promises for the sake of our promisees: in light of this new self-understanding, it would be morally impermissible for protagonists to make their promises in light of a difference-seeming. The hardliner, then, gets it right

both on the assumption that difference-seemings are epistemically significant and on the assumption that they're not.

Yet another thing we want out of an approach to the promising puzzle is that it avoid recommending epistemic evasion, or else explain why it's not morally problematic. The hardline response can. First, it's worth noting that the potency of the worry that hardline responses promote epistemic evasion is already diluted quite a bit – as we've seen, the hardliner can say that protagonists should promise in many puzzling cases. This means that the hardliner recommends the line of reasoning characteristic of epistemic evasion – “well, since I have evidence that there's a significant chance I won't follow through, then I just won't make the promise” – a lot less than Marušić contends.

And in the cases where the hardliner does recommend this line of reasoning – namely, in the puzzling cases just surveyed where the hardliner must say that protagonist shouldn't promise – it's exactly the line to take. When protagonists refuse to commit and promise to φ in these cases, they are not shirking responsibility for their future actions. Instead, constrained by self-awareness and caution, they're taking steps to lessen the damage of their future actions. That's not evading responsibility – that's taking it!

The hardline response can also account for the asymmetry between what the protagonist should believe and what a detached observer should believe about whether the protagonist will φ in puzzling cases – at least in those puzzling cases where the protagonists should promise. In these cases, it turns out, the protagonists' total body of evidence supports the belief in question while the detached observer's total body of evidence doesn't. That's because the protagonists – but not the detached observer – have access to the evidence that provides them with good reason to believe they are not relevantly like the comparison group.

In no-symmetry cases, this evidence largely comes from knowledge of the various desires, cares, values, beliefs, and other motivational and affective states that determine whether the protagonists will φ . Many such states are mental states, and, of course, protagonists have unique access to their own mental states. So protagonists alone – not the detached observed – have unique access to the evidence about

these features which provides them with good reason to believe that they're not relevantly like the comparison group.

In a good many broken symmetry cases, I've suggested this evidence comes from a difference-seeming, which is also a mental state, so the same thing can be said. What's more, only the possessor of a difference-seeming can acquire good reason to believe its propositional content – since a difference-seeming is had only by the protagonist, not by detached observers, only the protagonist reaps its epistemic benefits.

To conclude: the hardline response gives us everything we want out of an approach to the promising puzzle. Not only does it render the right verdicts in puzzling cases and concedes no ground on promising against the evidence, but it also avoids recommending epistemic evasion and makes room for the doxastic asymmetry between promiser and detached observer. This means we can keep our three original claims and also our promises to do things that many others, or we ourselves, have tried and failed to do, all without bearing the moral burdens of promising against the evidence.

Acknowledgements

This paper began as a term paper for a seminar on the ethics of belief at Harvard co-taught by Susanna Rinard and Selim Berker. It owes much to this seminar. For tremendously helpful feedback, I would like to thank Zach Barnett, Louis Gularte, Emily Hodges, Chad Marxen, Joshua Schechter, Nomy Arpaly, Jamie Dreier, the members of the Spring 2020 Dissertation Workshop at Brown University, and the anonymous reviewer at *Philosophers' Imprint*. Special thanks to David Christensen.

References

- Alicke, Mark. 1985. "Global Self-Evaluation as Determined by the Desirability and Controllability of Trait Adjectives." *Journal of Personality and Social Psychology* 49: 1621 – 1630.
- Downie, R.S. 1985. "Three Accounts of Promising." *Philosophical Quarterly* 35: 259–271.
- Friedrich, Daniel and Nicholas Southwood. 2011. "Promises and Trust." Hanch Sheinman (ed.), *Promises and Agreements* (ed. Hanoeh Sheinman): 277 – 294. Oxford: Oxford University Press.
- Fox, Richard and Joseph DeMarco. "On Making and Keeping Promises." *Journal of Applied Philosophy* 13: 199 - 208.
- Hoorens, Vera. 1993. Self-Enhancement and Superiority Biases in Social Comparison. *European Review of Social Psychology* 4: 113 – 139.
- Liberto, Hallie. 2016. "Promissory Obligation: Against a Unified Account." Mark Timmons (ed.), *Oxford Studies in Normative Ethics, Volume 6*: 102 – 122. Oxford: Oxford University Press.
- Marušić, Berislav. 2013. "Promising Against the Evidence." *Ethics* 123: 297 – 317.
- Marušić, Berislav. 2015. *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford: Oxford University Press.
- Lieberman, Alida. 2019. "Permissible Promise-Making Under Uncertainty." *Journal of the American Philosophical Association* 5: 468 – 486.
- Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94: 121-143.
- Rinard, Susanna. 2018. "Equal Treatment for Belief." *Philosophical Studies* 176: 1923 – 1950.
- Saul, Jennifer. 2003. *Feminism: Issues & Arguments*. Oxford: Oxford University Press.
- Scanlon, Thomas. 1990. "Promises and Practices." *Philosophy and Public Affairs* 19: 199 – 226
- Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Shiffrin, Seana. 2008. "Promising, Intimate Relationships and Conventionalism." *Philosophical Review* 177: 481–524.
- Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge: Harvard University Press.
- Vavova, Katia. 2018. "Berislav Marušić, *Evidence and Agency: Norms of Belief for Promising and Resolving*." *Ethics* 128: 687 – 695.