

## PROMISES, OFFERS, REQUESTS, AGREEMENTS

BRENDAN DE KENESSEY

*University of Toronto*

If I promise to pick you up at the airport, I thereby become obligated to do so. But this is not the only way I could undertake this obligation. If I *offer* to pick you up, and you accept my offer, I become obligated to pick you up in much the same way. I would also undertake similar obligations if you asked me to pick you up and I accepted your *request*, or if we made an *agreement* that I will pick you up at the airport and in exchange you'll buy me dinner. Why are the normative effects of accepted offers, accepted requests, and agreements so similar to those of promises? I argue that theorists of promising need to answer this question, and so they need to pay attention to offers, requests, and agreements. On the theory I defend, promises, offers, requests, and agreements have such similar normative effects because they all result in *joint decisions* between the relevant parties. I argue that this 'joint decision view' provides an attractive explanation of the similarities and differences between promises, offers, requests, and agreements.

**S**UPPOSE I promise you that I'll read a draft of yours by Friday. Under normal conditions, my promise will create an obligation for me to read your draft by Friday. I owe this obligation to you in particular, and you alone have the power to release me from it.

Now suppose instead I *offer* to read your draft by Friday, and you accept my offer. My offer, combined with your acceptance, has a notably similar pattern of normative effects. Because you accepted my offer, I am now obligated to read your draft by Friday; I owe this obligation to you in particular; and you alone have the power to release me from this obligation.

Now suppose you ask me to read your draft by Friday, and I accept your *request*. Your request, combined with my acceptance, seems to have the same normative effects. Because I accepted your request, I am now obligated to read

---

**Contact:** Brendan de Kenessey <brendan.dekenessey@utoronto.ca>

your draft by Friday; I owe this obligation to you in particular; and you alone have the power to release me from this obligation.

Finally, suppose we *agree* that I'll read your draft by this Friday and you'll read a draft of mine next week. Our agreement's effects on my normative situation fit the same pattern. I am now obligated to read your draft by Friday, and I owe this obligation to you. And, though you now need my assent to do so, you have the power to release me from our agreement.

This shared pattern of moral impact suggests that promises, offers, requests and agreements are intimately related. Yet philosophers have focused their attention almost exclusively on one member of this family: promises.<sup>1</sup> As a result, they have overlooked a pressing question: why are the normative effects of accepted offers, accepted requests, and agreements so similar to those of promises?

If our theory of promising is unable to answer this question, that is cause for worry. Imagine a theory of electricity that cannot be extended to magnetism, or a theory of the rise of Nazi Germany that has nothing to say about Stalinist Russia or Maoist China. The problem with these theories is not merely that they have neglected interesting topics also worthy of study. Instead, each theory's inability to explain nearby phenomena indicates that it has missed something important about its own subject: the laws of electromagnetism or the origins of totalitarianism.

This paper begins by arguing that theories of promising face an analogous problem if they neglect offers, requests, and agreements. Section 1 details the similarities between the normative effects of promises and those of accepted offers, requests, and agreements. These similarities strongly suggest that promises, offers, requests and agreements have a common basis. So, if our theory of promising cannot explain the similar effects of offers, requests, and agreements, this gives us reason to doubt its adequacy *as a theory of promising*.<sup>2</sup>

An easy answer to this challenge suggests itself. Perhaps offers, requests, and agreements are so similar to promises because they *are* promises—or at least involve them. On this view, my offering to read your draft, accepting your request to read it, or agreeing to read it are just different ways of my *promising* to read your draft. Call this hypothesis *promissory reductionism*. If promissory reductionism is right, our challenge can be quickly dismissed. Any theory of promises can explain offers, requests, and agreements, as these phenomena are all just promises in disguise.

---

1. There has been some discussion of agreements (Gilbert 1993; 2018; Bach 1995; Sánchez Brígido 2006; Sheinman 2011; Black 2012) and requests (Raz 1986; Cupit 1994; Enoch 2011; Hanser 2015; Gläser 2019). The only philosophical discussion of offers I have found is about whether offers can be coercive (Häyry & Airaksinen 1998; Rocha 2011).

2. Fruh (2014; 2019) pursues a similar strategy with different materials, arguing that theories of promising should be able to capture the similar phenomena of oaths, vows, and promises to oneself. I am not sure what to think about oaths and vows, but I address promises to oneself in §4.4.

Section 2 argues that promissory reductionism is not obviously true. Considering the most straightforward ways of reducing offers, requests, and agreements to promises, I argue that each reduction either does not work or leaves important questions unanswered. My aim is not to refute promissory reductionism, but rather to show that it is not a quick and easy solution to our problem.

Section 3 proposes an alternative way of unifying our four phenomena. Some philosophers, including myself, have recently defended the claim that promises should be understood as resulting in joint decisions: when I promise that I'll read your draft, and you accept my promise, the result is a joint decision between us to the effect that you'll read my draft (Gilbert 2018; de Kenessey 2020). I argue that this *joint decision view* can be naturally extended to offers, requests, and agreements. On the theory I propose, promises, offers, requests, and agreements are four different ways of deciding together what to do. Promises and offers propose joint decisions concerning what *I* will do; requests propose joint decisions concerning what *you* will do; and agreements are joint decisions concerning what *we* will do. The similar normative effects of these four acts are just the normative effects of making a joint decision. In turn, the differences between these four phenomena can be explained by appeal to the different ways in which they propose joint decisions.

Section 4 concludes by considering objections.

Most philosophy papers aim to answer old questions; some set out to ask new ones. This paper is of the latter kind. My primary aim in this paper is to draw my colleagues' attention to the similarities between promises, offers, requests, and agreements, and pique their curiosity about how these similarities are to be explained. Since I am trying to start a conversation rather than end it, this paper's positive ambitions are accordingly modest. Neither my criticisms of promissory reductionism nor my defense of the joint decision view are intended to be conclusive. Like Tom Sawyer at the picket fence, my goal is not to finish the job myself, but to make the work look interesting enough to tempt you to join in.

## **1. The Similarities between Promises, Offers, Requests, and Agreements**

### ***1.1. Four Examples***

This section argues that the obligations that result from accepted offers, accepted requests, and agreements are similar to those generated by promises. To begin, consider

**PROMISE:** *You and I are friends from graduate school, now living in different cities. One day you call me and say:*

You: Hey, guess what? I'm coming to your university next week to give a talk. I'm flying in on Thursday.

ME: That's great news! I'll pick you up from the airport.

You: Great, thanks so much. Can't wait to see you!

By saying, "I'll pick you up from the airport," I have promised you that I will do just this. This promise creates a new obligation for me to pick you up from the airport.

Now consider

**OFFER:** *You and I are friends from graduate school, now living in different cities. One day you call me and say:*

You: Hey, guess what? I'm coming to your university next week to give a talk. I'm flying in on Thursday.

ME: That's great news! Hey, would you like for me to pick you up from the airport?

You: That would be great, thanks so much. Can't wait to see you!

**REQUEST:** . . . *you call me and say:*

You: Hey, guess what? I'm coming to your university next week to give a talk. I'm flying in on Thursday. Actually, I was wondering: would you be willing to pick me up from the airport?

ME: Sure, I'd be happy to do that.

You: That's great, thanks so much. Can't wait to see you!

**AGREEMENT:** . . . *you call me and say:*

You: Hey, guess what? I'm coming to your university next week to give a talk. I'm flying in on Thursday. How about you pick me up from the airport, and in return I'll buy you dinner on the way back?

ME: Sure, that sounds good.

You: Great, thanks so much. Can't wait to see you!

Here is my claim: the obligation to pick you up from the airport that results from each of these conversations is identical to the obligation that results from PROMISE. The only normative difference between these interactions is that in AGREEMENT, you also undertake an obligation to buy me dinner. Otherwise, PROMISE, OFFER, REQUEST, and AGREEMENT appear to have precisely the same result.

Here's an argument for this claim. Suppose that, by Thursday morning, I have forgotten exactly what was said when we spoke on the phone last week. I

don't remember whether our interaction was as depicted in PROMISE, or instead more like OFFER, REQUEST, or AGREEMENT. Should my uncertainty about which conversation we had make me uncertain about the nature of my obligation to pick you up from the airport? Intuitively, the answer is no. It doesn't matter which conversation we had. Whichever of the above dialogues transpired, my obligation is the same. Again, the only difference is that in AGREEMENT, you are obligated to buy me dinner. But as far as my obligation to you is concerned, it does not seem to make a difference whether we arrived at it via PROMISE, OFFER, REQUEST, or AGREEMENT. This strongly suggests that promises, accepted offers, accepted requests, and agreements result in the same normative effects.

### 1.2. A Shared Pattern of Normative Features

Here I bolster the above argument by reviewing several features of promissory obligation and arguing that the same features apply to the obligations generated by accepted offers, accepted requests, and agreements. Not only do accepted offers, requests, and agreements share promises' obligation-generating power, but the *pattern* of the obligations they generate is remarkably similar to that of promissory obligations.

(1) *Obligation*. Under normal conditions, a promise to  $\phi$  places the promisor under a moral obligation to  $\phi$ . The obligations created by promises are strong but not absolute. If my child falls severely ill and needs to be taken to the hospital on Thursday, then it is intuitively permissible for me to break my promise to pick you up. Even when it is permissible, however, promise-breaking places the promisor under a secondary duty to 'make up' for the broken promise in some way: for example, by paying for the taxi you had to hire instead.

Similarly, under normal conditions, when you accept my offer to  $\phi$ , or I accept your request that I  $\phi$ , or we agree that I will  $\phi$ , this places me under a moral obligation to  $\phi$ . The obligations created by accepted offers, accepted requests, and agreements are strong but not absolute. But, as with promises, even if I am justified in failing to fulfill an accepted offer, accepted request, or agreement, my failure places me under a secondary duty. If I strand you at the airport after OFFER, REQUEST, or AGREEMENT, then I am obligated to pay for your taxi just as I would be after PROMISE.

Some might balk at the claim that offers and requests generate obligations. And rightly so: offers and requests do not generate obligations *on their own*, but only when combined with the other party's acceptance. Start with offers: the fact that I offered to pick you up is not enough to obligate me to do so, for you might ignore or turn down my offer. But if I offer to pick you up and you accept my offer, then I am (normally) obligated to do so. Strictly speaking, then, it is not

offers that generate obligations, but *completed offering transactions*, composed of the making and accepting of an offer. I will call these ‘accepted offers’ for short.

The same holds for requests: the fact that you requested that I pick you up is not enough on its own to obligate me to pick you up. Intuitively, it is part of the very idea of a request that the addressee is able to turn it down. For this reason, philosophers have tended to deny that requests generate obligations (e.g., Raz 1986: 37). These philosophers are right that requests do not generate obligations on their own. But a *completed requesting transaction*, composed of the making and accepting of a request, does generate an obligation. I will call these ‘accepted requests’ for short.

(2) *Deontological structure*. Promissory obligations are widely taken to be a paradigm case of a deontological constraint. This means that it can be the case that I ought to keep my promise even if breaking it would produce a better state of affairs. Even if I could produce more goodness by spending Thursday volunteering, I intuitively still ought to pick you up. The fact that I promised gives me a non-instrumental reason to pick you up that does not depend on its promoting any other valuable outcome. Moreover, the reason to keep a promise is agent-relative: I have reason to keep *my* promises, not to see to it that promises in general are kept.

The obligations created by accepted offers, accepted requests, and agreements share this deontological structure. Like the reason to keep a promise, the reason to fulfill an accepted offer, accepted request, or agreement is non-instrumental and agent-relative. This reason is independent from any further value that might result from the action; and it is a reason to do what *I* offered (or accepted a request, or agreed) to do, not to see to it that accepted offers, accepted requests, or agreements are fulfilled in general.

(3) *Directedness*. It is often observed that promissory obligations are owed *to* the promisee (e.g., Gilbert 2018). When I promise to pick you up, I become obligated to you in particular to do so. If I don’t pick you up, I will have wronged *you*. The obligations created by offers, requests, and agreements are similarly directed. If I fail to follow through on my offer that you accepted, your request that I accepted, or an agreement we made, I wrong you in particular just as surely as if I broke a promise to you.

(4) *Waivability*. Promisees have the power to *release* promisors from their promises, thereby waiving their promissory obligations. If you call me up and say “Actually, don’t pick me up—I decided to rent a car instead,” you have thereby released me from my obligation. This power of release is held by the promisee alone: no one else can waive a promissory obligation.

Similarly, you alone have the power to release me from the obligations I undertake by making an offer that you accept, accepting your request, or making an agreement with you. You could just as easily call me after our conversation in OFFER, REQUEST, or AGREEMENT and waive my obligation to pick you up.



For agreements in particular, however, it is less apt to describe one party as having ‘the power of release’. Since agreements generate mutual obligations, the power to waive those obligations is similarly symmetrical. Instead of one of us releasing the other, we have to jointly release ourselves. Rather than simply saying “You don’t need to pick me up,” you might say something like: “Actually, do you mind if I take a cab to my hotel, and we get dinner another time?” If I say, “Sure, that’s fine,” we have released ourselves from our agreement.

(5) *Exclusionary force*. Some philosophers have argued that promises have an *exclusionary* effect on deliberation (e.g., Raz 1977). Once you have promised to do something, you should treat it as already decided that you will do it, and should exclude at least some reasons for breaking your promise from your deliberation. Suppose that on Thursday I deliberate as follows: “What should I do this afternoon? I know I promised to be at the airport, but I’d rather just watch TV . . .” Even if I ultimately keep my promise, there is something off about this thought process: I should not be thinking of breaking my promise as a live option at all.

This claim appears equally plausible for accepted offers, requests, and agreements. After OFFER, REQUEST, or AGREEMENT, it would be equally inappropriate for me to deliberate about whether to leave you stranded at the airport so as to watch TV.

(6) *Invalidation by coercion and deception*. Promises lose their obligating force when they are elicited by coercion or deception. If you hold me at gunpoint and tell me to promise that I will send you \$10,000 when I get home, and I prudently promise just that, I am intuitively under no obligation to keep this promise. If you get me to promise to buy your painting by falsely telling me that it is an undiscovered Warhol, then once I learn of your deception, I am intuitively under no obligation to buy your painting. (These examples are from Thomson 1990: 310–13.)

Accepted offers, accepted requests, and agreements are invalidated by coercion and deception as well. To see this, just tweak the above cases. Suppose you hold me at gunpoint and tell me to *offer* to send you \$10,000, or you *ask* me to send you \$10,000, or you propose that we *agree* that I’ll send you \$10,000 and you’ll let me go. Or suppose that, after you falsely tell me you have a Warhol, I *offer* to buy it, and you accept; or you *ask* me to buy your Warhol, and I accept your request; or we *agree* that I’ll pay you \$10,000 and in exchange you’ll give me your Warhol. None of these transactions obligate me any more than the parallel promises do.

### 1.3. Why Theories of Promising Need an Account of Offers, Requests, and Agreements

Promissory obligations have a characteristic, fingerprint-like pattern of normative features, one that philosophers have mapped in great detail. When we compare

the obligations generated by accepted offers, accepted requests, and agreements, they seem to replicate this pattern exactly. The match is too robust to be coincidental: the shared normative pattern of promises, offers, requests, and agreements calls out for explanation.

I believe that theories of promising should be expected to answer this call. We should expect our theory of promising to explain why the obligations generated by accepted offers, accepted requests, and agreements are so similar to those generated by promises. If a theory of promising cannot do this, that is a count against it.

Some theorists of promising may chafe at this claim. They may protest that the subject-matter they set out to explain is *promising*. Explaining the normative effects of accepted offers, accepted requests, and agreements is a valuable theoretical project, they may say, but it is not *their* project. Their project is to give a theory of promises, and that is hard enough as it is. Philosophy requires some division of labor, after all!

While philosophy does require some division of labor, it is important to avoid dividing our labor in a way that blinds us to theoretically important considerations. And it seems clear to me that the fact that accepted offers, accepted requests, and agreements generate promise-like obligations is a consideration that is relevant to our theory of promises.

One argument for this claim appeals to the idea that similar phenomena are likely to have similar explanations. To illustrate, suppose an eighteenth-century chemist proposes a theory of why water freezes at low temperatures. While the theory predicts all of our observations about water, it has nothing to say about why *other* liquids, such as mercury or oil, freeze at low temperatures. This failure to explain similar phenomena would be a strike against the theory. The freezing of mercury and oil is similar enough to the freezing of water that, in the absence of evidence to the contrary, we should expect these processes to have a common explanation. The fact that our chemist's proposed theory cannot explain why other liquids freeze indicates that it is missing something important about why *water* freezes.

Analogously, the obligations produced by promises, accepted offers, accepted requests, and agreements are similar enough that we should expect them to have a common explanation. Recall the examples in §1.1: PROMISE, OFFER, REQUEST, and AGREEMENT look like four slightly different routes to the same normative result. Given their close similarity, it would be surprising if the obligations resulting from these four conversations had different explanations: if, say, the obligation to pick you up from the airport I undertake in PROMISE had a fundamentally different basis from the obligation I undertake in OFFER, REQUEST, and AGREEMENT. The presumption should be that these matching normative effects have a common normative 'cause'.

So, imagine a theory of promises that cannot be extended to explain offers, requests, and agreements. Suppose that, for whatever reason, the obligations



generated by accepted offers, accepted requests, and agreements cannot be adequately understood in the way the theory explains promissory obligations. As in our chemistry example, this failure to explain similar phenomena would be a strike against the theory. The normative effects of accepted offers, accepted requests, and agreements are similar enough to the normative effects of promises that we should expect them to have a common explanation. And so, the fact that our hypothetical theory cannot explain why accepted offers, accepted requests, and agreements generate obligations indicates that it is missing something important about why promises generate obligations.

A second argument for the same conclusion appeals to the *purpose* of studying promising. Why are philosophers interested in promises in the first place? I take it that philosophers' primary aim in studying promises is to understand promises' power to create obligations where none existed before. After all, the modern literature on promises was kicked off by Hume's declaration that promises' ability to conjure obligations is "one of the most mysterious and incomprehensible operations that can possibly be imagin'd" (1739/1978: 524). The aim of a theory of promises is to demystify this operation, explaining how promises make it possible for us to create new obligations by declaration. But as we have just seen, accepted offers, accepted requests, and agreements also create obligations where none existed before, performing acts of normative transubstantiation that Hume would have found no less mysterious. So if the philosophical purpose of studying promises is to understand how we can create new obligations by declaration, that purpose requires us to also ask how accepted offers, accepted requests, and agreements accomplish this feat.

I conclude that theorists of promising cannot simply declare that offers, requests, and agreements are somebody else's problem. The question "How do promises create obligations?" is too narrow in the same way that "Why does water freeze at low temperatures?" is. If the chemist focuses solely on water and does not consider why other liquids also freeze at low temperatures, she will overlook the fact that the phenomenon she is trying to explain—how liquids turn into solids—is more general, and thus requires a more general theory. Similarly, if philosophers focus solely on promises and do not consider why accepted offers, accepted requests, and agreements also create obligations, they will fail to see that the phenomenon they are trying to explain—how people can undertake new obligations by declaration—is more general, and thus requires a more general theory.

## **2. Promissory Reductionism: An Easy Solution?**

Suppose I've convinced you that our theory of promising needs an explanation of why accepted offers, accepted requests, and agreements generate promise-like

obligations. You may still think that addressing this challenge won't require you to revise your theory of promising, or really to do much work at all. For you might think that there is an easy explanation: accepted offers, accepted requests, and agreements have such similar normative effects to promises because they *are* promises, or involve them somehow.

Let us call *promissory reductionism* the view that the normative effects of accepted offers, accepted requests, and agreements can be exhaustively explained in terms of promises. On this view, there is really only one obligation-conjuring trick in our normative spellbook: promising. OFFER, REQUEST, and AGREEMENT all result in a promise-like obligation for me to pick you up from the airport because, at some point in each of these conversations, I *promise* you that I will pick you up from the airport.

If promissory reductionism is true, it will be easy for theorists of promising to explain offers, requests, and agreements. If these acts are promises in disguise, a theory of why promises generate obligations will automatically also explain the obligations generated by accepted offers, accepted requests, and agreements.

Establishing the truth of promissory reductionism itself, however, is less easy. This section argues that the most straightforward attempts to reduce accepted offers, accepted requests, and agreements to promises each face substantial obstacles. My aim in highlighting these obstacles is not to demonstrate that promissory reductionism is false, but just to show that it faces *prima facie* challenges that will take real work to overcome. My wider goal in this paper is to motivate the need for a debate about the relation between promises, offers, requests, and agreements—not to settle that debate or rule out any contenders. And so my modest hope in this section is to show that promissory reductionism is not so clearly correct as to end the debate before it starts.

### 2.1. Offers

The most salient difference between promises and offers concerns their requirements for uptake. Offers do not result in obligations unless and until they are explicitly accepted by their addressees. In OFFER, when I say “Would you like for me to pick you up from the airport?” I am not yet obligated to pick you up from the airport. This obligation only comes into force when you accept my offer, saying “That would be great, thanks.” In contrast, while promises do require uptake from their addressees, they do not require an overt act of acceptance.<sup>3</sup> Instead, all

---

3. This claim appears to put me in conflict with philosophers who argue that promises must be accepted by the promisee (Fried 1981: 41–43; Thomson 1990: 296–98; Owens 2006: 72; Roth 2016: 89–92; Gilbert 2018: 108–10). For Fried, Thomson, and Owens, the disagreement is only apparent, as these theorists allow that a promise is tacitly accepted if the promisee understands it and does

a promise requires for uptake is that the promisee understands the promise and does not object to it. And so, in PROMISE, I become obligated the moment I say “I’ll pick you up from the airport.”

This difference presents a challenge for the promissory reductionist. To reduce offers to promises, she must identify an act of promising somewhere in a completed offering transaction. The natural candidate is the offer itself: when I say, “Would you like me to  $\phi$ ?” I am promising, in some way, to  $\phi$ . The trouble is that, unlike a promise, an offer is not enough on its own to generate an obligation. Instead, the speech act that directly generates the obligation appears to be the addressee’s acceptance of the offer. But the addressee’s acceptance cannot be the promise that binds the offeror’s actions—unless we want to allow for promises that bind other people’s actions, as in the mafioso-like “I promise that *you* will pick me up from the airport.” The question becomes: if offers are promises, why is there a delay between the offer and the creation of an obligation?

Here’s a simple answer: perhaps offers require explicit acceptance because they are *conditional* promises. When I offer to pick you up from the airport, it is as if I had said, “I promise you this: if you say you want me to pick you up, I will do so.” And so I do not become obligated to pick you up until you say you want me to do so, as it is built into the content of my promise that this obligation is conditional on your acceptance.

This hypothesis solves our first problem, but runs into a second: it fails to capture offers’ *revocability*. Suppose I offer to give you my antique bookcase. The bookcase is nice, but you’re not sure you have space for it. “I’ll think about it,” you say. Weeks pass, and then one day, someone else offers to buy the bookcase at a good price. I can then call you up and say, “Sorry, something came up—my bookcase isn’t available anymore.” I am perfectly within my rights to do this. This shows that offers are revocable: until you accept my offer, I have the power to revoke it.<sup>4</sup>

The trouble is that promises *cannot* be revoked unilaterally. I cannot simply take back a promise when I don’t want to fulfill it anymore. And this is just as true for conditional promises. Suppose you are playing the Pac-Man machine at our favorite bar. I say, “I promise you this: if you beat my high score, I will buy you a beer.” If I see you racking up points, I can’t then say, “Never mind, I take it back.”

If an unaccepted offer is just a conditional promise whose condition has yet to be fulfilled, then it should be similarly impossible for this promise to be revoked. The fact that unaccepted offers *can* be revoked is thus evidence that they are not simply promises conditioned on their addressees’ acceptance.

---

not object. However, both Roth and Gilbert make the stronger claim that promises must be overtly accepted to receive uptake (though Roth is ambivalent—see 2016: 92, fn. 9). I disagree.

4. It is standard doctrine in contract law that offers can be unilaterally revoked before they are accepted (Benson 2019: 102–10).

I see two potential responses to this problem. The first response is to build an implicit revocability clause into the conditional promise. When I offer you my bookcase, perhaps I am not making the simple promise “If you tell me you want it, I’ll give it to you,” but rather a more complex promise like: “If you tell me you want it *before I tell you it is no longer available*, then I’ll give it to you.”

This hypothesis implies that it is possible to build a revocability condition into promises—indeed, that we do this tacitly every time we make an offer. This would lead one to expect to find some promises where the revocability condition is made explicit. But such promises are nowhere to be seen. Imagine if I said: “I promise you this: if you beat my high score on Pac-Man, I will buy you a beer, unless I say otherwise before you do.” That’s not a promise—that’s a joke. To say “I promise, unless I say otherwise” is to undercut my promise so dramatically that it is hard to see it as a promise at all (cf. Shiffrin 2007: 728–29). Moreover, most theorists agree that the purpose of making a promise is to undertake a commitment that isn’t easy to escape. This purpose appears to be undermined if we allow promisors to build in an ‘opt-out’ clause. The proposed reduction implies not only that such opt-out clauses are possible, but that they are pervasive—not an attractive consequence.

A second response available to the promissory reductionist is to maintain that the appearance that offers can be revoked is explained by the fact that they can be *overridden*. Promises generate only pro tanto obligations, which can be permissible to infringe if one has overriding reason to do so. The reductionist might propose that this is what is going on in cases where it seems permissible to revoke an offer. The reason why I can revoke my offer when someone offers me money for my bookcase is that this gives me sufficient reason to break my conditional promise to give you the bookcase if you want it. And so there is no disanalogy between offers and conditional promises after all—both can be ‘revoked’ when and only when there is sufficient reason to break them.<sup>5</sup>

The problem with this response is that there *is* a disanalogy between offers and conditional promises. If I had promised you “If you ask for the bookcase, I will give it to you,” then I would need quite a strong reason to justify breaking this promise. Similarly, if you had accepted my offer to give you the bookcase, I would need a strong reason to justify not giving it to you. But before you accept my offer, I can revoke it for a much weaker reason. I would be within my rights to take my offer back simply because I changed my mind, deciding I wanted to keep the bookcase after all. Even if I need *some* reason to justify revoking my offer, this reason is surely weaker than the reason I would need to break a promise.

Remember that the promissory reductionist’s hypothesis is that to offer to give you the bookcase *just is* to promise to give it to you if you ask for it. And so

---

5. Thanks to a referee for suggesting this reply.

their view leaves no room for differences between these acts. The ‘conditional promise’ hypothesis cannot explain the fact that it is harder to justify breaking a conditional promise than it is to justify revoking an unaccepted offer.

There may be other responses available to the promissory reductionist. But the arguments so far are enough to make my point: offers are not *easily* reducible to promises.

## 2.2. Requests

To reduce requests to promises, the promissory reductionist needs to identify an act of promising somewhere within a completed requesting transaction. Since promises obligate the person who makes them, we cannot identify the promise with the request itself: for it is not the person making the request, but the person accepting it, who undertakes an obligation. This suggests a reductionist hypothesis: perhaps to accept a request *just is* to promise to fulfill it. When I accept your request to pick you up from the airport, I am promising to do just that.

This ‘accepting as promising’ hypothesis delivers all the right predictions about what happens after a request is accepted—for, as I have contended, accepted requests generate promise-like obligations. My worry concerns the implications of this hypothesis for our understanding of requests themselves. If to accept a request is to make a promise, then what is it to *make* a request? What is the relation between the act of making a request and the act of accepting one?

Intuitively, the acts of making and accepting a request are two parts of a single normative transaction. Compare the acts of giving and accepting a gift, or of proposing and accepting marriage, or for that matter, the acts of making and accepting an offer. In these cases, the act of accepting (a gift, a marriage proposal, or an offer) completes a process that was initiated by the other party’s action (of giving, proposing, or offering). As a result, the accepting act is not intelligible independently from the initiating act: for example, one cannot accept a gift that was not first given. I think the same is true of requests: the act of accepting a request is not intelligible independently from the act of making one.

The accepting-as-promising hypothesis paints a different picture. On this view, to accept a request is to make a promise. But it is not only intelligible, but commonplace, to make a promise that does not respond to a request. And so accepting-as-promising implies that the act of accepting a request is separable from the act of making one. While we would not describe a promise as ‘accepting a request’ unless it was made in response to a request, the normative essence of promising is the same whether or not it is elicited by a request. Rather than being two parts of a single transaction, making and accepting requests become two normatively independent acts. The making of the request merely presents

an occasion for the addressee to make a promise, one that she just as easily could have made without it. The addressee's promise does all the heavy lifting in generating her obligation: the request itself appears to be normatively inert.

I find this picture of the relation between making and accepting requests counterintuitive, but others may not share my intuition. So I'll phrase my objection as a challenge. To give a complete account of requests, the promissory reductionist needs to offer us an account of what it is to make a request. The challenge is for this account to explain why requests are so tightly connected to promises: what makes it the case that to accept a request is to promise to fulfill it?

I'll consider two proposals on the promissory reductionist's behalf. These do not exhaust the reductionist's options, of course; the point is to show that the task is non-trivial.

The first proposal is that to request is to *invite* the addressee to promise to do the requested action.<sup>6</sup> This proposal nicely explains the connection between requests and promises. If to request is to invite a promise, then to accept a request would be to accept that invitation, and thereby to make the invited promise. My worry about this proposal is that the concept of an invitation is too close to that of a request for this analysis to be informative. I have a hard time seeing much conceptual distance between requests and invitations. Both acts propose an action to their addressee; both are such that, if the addressee accepts them, they take on a promise-like commitment. This leads me to suspect that, despite the different connotations of the English words 'invite' and 'request', there is no deep distinction between the speech acts these words denote.

If I squint a bit, I can discern a concept of invitation that is distinct from that of a request. But on this conception, it seems possible to invite without making a request. I might say to a socially anxious friend, "You're invited to my party, but I'm not asking you to come—it's totally up to you." This gives my friend permission to attend the party without asking them to attend. The same holds for invitations to promise: suppose I invite my friend to RSVP to the party, making it clear that I will treat a 'yes' RSVP as a promise to come. I might still say: "But I'm not asking you to come—it's totally up to you." Insofar as I have an independent grip on the concept of an invitation, it seems to have a different extension from requests. So, while the request-as-invitation view secures the right connection between requests and promises, it appears to founder on a different problem.

Our second proposal comes from Joseph Raz: requests aim to influence their addressees by *giving them a reason*. As Raz puts it: "Requests . . . are acts intended to communicate to their addressee the speaker's intention that the addressee shall regard the act of communication as a reason for a certain action" (1986: 37; Cupit 1994 and Enoch 2011 also endorse this view). The promissory reductionist

---

6. Thanks to a referee for suggesting this view.



can coopt Raz's view by changing just one detail: they should say that a request to  $\phi$  gives its addressee reason to either  $\phi$  directly *or* promise to  $\phi$ .

I have two worries about this proposal. One is that there are independent objections to the reason-giving view. Micha Gläser has recently argued, persuasively in my view, that the reason-giving view yields overly permissive predictions about when it is appropriate to make a request (Gläser 2019). If the promissory reductionist wishes to adopt Raz's reason-giving view, they will need to respond to Gläser's objections.

My second worry is that the reason-giving view does not deliver the connection between requesting and promising that we are seeking. On the reason-giving view, requests just add one pro tanto reason to the balance of considerations the addressee should weigh up when deciding what to do. As Raz emphasizes: "The speaker . . . intends to influence [the addressee] only by tipping the balance somewhat in favour of the requested act" (1986: 37). It seems to follow that, on the reason-giving view, to accept a request would be to accept the reason it offers. But I can accept a pro tanto reason to  $\phi$  without promising to  $\phi$ . On the reason-giving view, then, it seems that one could accept a request by simply saying, "Thanks, I'll take that reason into consideration," while making no promise to fulfill it. This is still the case if we emphasize that requests provide reasons *to promise*—I can accept a pro tanto reason to promise without making the promise it supports. And so the reason-giving view does not explain why accepting a request involves promising to fulfill it, as the promissory reductionist must maintain.

I admit that the promissory reductionist is in a stronger position with regard to requests than they are with offers or—as we are about to see—agreements. The view that accepting a request is promising to fulfill it does not face any direct counterexamples as far as I can see. My worry is that this view does not, at least without further elaboration, adequately capture the connection between the acts of making and accepting requests. And so, again, my contention is not that promissory reductionism is false, but that it needs more defense.

### 2.3. *Agreements*

The most salient difference between agreements and promises is that promises bind only the speaker, while agreements generate obligations for all involved. When you and I agree to meet for coffee, we each become obligated to meet the other for coffee. The straightforward reductionist explanation of this is that both of us have made promises: I've promised you that I will meet you for coffee, and in exchange you've promised that you will meet me. Generalizing, our reductionist hypothesis is: for S and A to make an agreement that S will  $\phi$  and A

will  $\psi$  just is for S to promise A that [. . .] and for A to promise S that [. . .]. The ellipses in this formulation are meant to allow different views about the content of the exchanged promises. It could be simply that S promises A that S will  $\phi$  and A promises S that A will  $\psi$ ; or perhaps agreements involve promises with more complex contents.

In her classic paper “Is an Agreement an Exchange of Promises?” Margaret Gilbert argues that no exchange of promises can replicate the normative impact of an agreement (Gilbert 1993; see also Sánchez Brigido 2006 and Black 2012). Her arguments raise a significant enough problem for the promissory reductionist that they are worth reviewing here.

Gilbert considers the following agreement: “Peter: ‘Why don’t you walk Fido, and I’ll groom Tibbles?’ Rita: ‘Fine!’” (1993: 627). How might the reductionist reconstruct Peter and Rita’s agreement as an exchange of promises?

Begin with the simplest proposal: Peter promises to groom Tibbles, and Rita promises to walk Fido (1993: 635). The problem with this exchange is that it fails to capture the *interdependence* of the obligations generated by agreements (1993: 630). Suppose that, after Peter and Rita make their agreement, Peter tells Rita that he’s decided not to groom Tibbles after all. Intuitively, once Peter has failed to hold up his end of the agreement, Rita no longer owes it to Peter to hold up hers. In general: if one party to an agreement fails to do what the agreement requires of him, then the other party is no longer obligated by the agreement to do what it requires of her.

In contrast, pairs of promises are not interdependent. Say I promise you that I will read your draft by Friday, and separately, you promise to look after my cat next week. I fail to read your draft. Are you off the hook for cat-sitting? Clearly not. While agreements generate interdependent obligations, straightforward promise-exchanges do not.

The obvious solution is to build interdependence into the content of the exchanged promises. Suppose Peter and Rita exchange the following promises: Peter promises Rita that he will groom Tibbles *if* she walks Fido, and Rita promises Peter that she will walk Fido *if* he grooms Tibbles (1993: 640). This proposal succeeds in making Peter and Rita’s obligations interdependent: if Peter does not groom Tibbles, Rita is not obligated to walk Fido, and vice versa. But in doing so, it fails to generate the very obligations we were trying to explain. One way for Peter and Rita to keep their promises would be for Peter to groom Tibbles and Rita to walk Fido. But another, equally legitimate way for Peter and Rita to keep their promises is for both of them to do nothing. So long as Peter does not groom Tibbles, Rita is under no obligation to walk Fido, and vice versa.

Agreements are different: they do not treat joint inaction as normatively equivalent to joint fulfillment. Peter and Rita cannot fulfill their agreement by doing nothing; it will only be honored if Peter grooms Tibbles and Rita walks Fido.

Thus the promissory reductionist must either fail to capture the fact that agreements generate interdependent obligations, or fail to explain how they generate obligations at all. The upshot is that no exchange of promises can do what an agreement does: generate obligations that are both interdependent and reason-giving. I take this to be the most daunting problem faced by the promissory reductionist.

As I have emphasized, the arguments in this section have aimed only to shift the burden of proof. Many readers may be tempted to treat promissory reductionism as the default view, assuming that offers, requests, and agreements can be reduced to promises in some way or other. This assumption is too complacent. Promissory reductionism needs more defense if it is to provide an adequate answer to the question raised in Section 1. In the meantime, it is reasonable to consider alternative ways of explaining the unity of promises, offers, requests, and agreements. The next section presents the theory I favor.

### 3. The Joint Decision View

#### 3.1. *Introducing the View*

A new theory of promises has emerged in the past two decades, which I will call *the joint decision view*. According to this view, the result of a successful promise is the making of a joint decision between the promisor and promisee to the effect that the promised action will be performed. When I promise to pick you up at the airport and you accept my promise, then, we have jointly decided that I will pick you up at the airport. This joint decision grounds my promissory obligation.

The joint decision view's first and most prominent defender is Margaret Gilbert (2018). Gilbert argues that we can only capture promises' distinctive obligating force by appeal to a joint commitment between the promisor and promisee. More recently, I have defended the joint decision view on different grounds, arguing that it enables us to explain promises' "quirks and qualities" by deriving them from the structure of joint decision-making (de Kenessey 2020: 206).<sup>7</sup>

Here I hope to offer a new point in favor of the joint decision view: this view, I will argue, is well-placed to provide a unified account of promises, offers, requests, and agreements. If we expand the joint decision view of promises to encompass offers, requests, and agreements, the result is an explanatory account of both the similarities and differences between these four phenomena.

A starting premise of the joint decision view is that it is possible to make joint decisions concerning only one party's actions. Paradigmatically, joint decisions

---

7. See Lichter (2021) for criticism of the joint decision view.

concern both parties' actions, as when we decide that *we* will go for a walk. But promises, offers, and requests concern the actions of only one person—and so if we are to understand them in terms of joint decisions, these joint decisions must concern one person's actions. On reflection, such joint decisions seem possible. Planning a cookout, we might jointly decide that *you* will tend the grill, leaving it open what I will do. So, joint decisions are not limited to the content *I will  $\phi$  and you will  $\psi$* —we can also jointly decide that *I will  $\phi$*  or that *you will  $\psi$* .

With this assumption in place, we can expand the joint decision view to include offers, requests, and agreements. On this view, promises and accepted offers result in joint decisions concerning what the speaker will do. Requests, when accepted, result in joint decisions concerning what the addressee will do. Agreements result in joint decisions that concern both parties' actions.

To unpack this view in more detail, we need to look at the process of joint decision-making. So let's ask: how do two or more agents go about making a joint decision?

The obvious method goes something like this. First, one of the participating agents—let's say me—proposes a joint decision. Then the other agent—let's say you—considers the proposal, deciding whether to accept it. If I propose a joint decision and you explicitly accept it, under conditions of common knowledge, then *voilà!*—we have made a joint decision. Call this method of joint decision-making *the propose-and-ratify method*.

Propose-and-ratify is what first comes to mind when one thinks of joint decision-making. I say, "Shall we go on a walk?" You reply, "Sure!" And a joint decision is made. But this is not the only way we can go about making joint decisions.

By adopting the propose-and-ratify method, we make a subtle but substantive choice of deliberative procedure: we assume that the default result is for a proposed joint decision to be *rejected*. On propose-and-ratify, if I propose a joint decision and you say nothing in response, then no joint decision is made. But it seems possible to adopt a deliberative method on which the default result is for a proposal to be *accepted*, and overt action must be taken to reject it.

I call this the *propose-and-challenge method*. To make a joint decision via propose-and-challenge, we start again with one party—say, me—proposing a joint decision I take to be justified. But then, instead of accepting the proposal if you take it to be justified, you are expected to *challenge* my proposal if you take it to be *unjustified*. If you recognize my proposal and refrain from challenging it, then it counts as accepted by default, and we have made a joint decision. Thus, when we are using propose-and-challenge, we count as having made a joint decision if I propose a decision and you say nothing in response. Assuming you understand my proposal, your silence counts as acceptance.

If we distinguish between the propose-and-challenge and propose-and-ratify methods, and allow that these methods may be used to reach joint decisions

concerning what *you* or *I* will do alone as well as what *we* will do, we get a set of categories into which promises, offers, requests, and agreements fit naturally:

	Propose-and-challenge	Propose-and-ratify
"I will $\phi$ "	Promises	Offers
"You will $\phi$ "	?	Requests
"We will $\phi$ "	Agreements	

On this view, promises and offers are both proposals to make joint decisions concerning the *speaker's* actions ("I will  $\phi$ "). The difference between promises and offers lies in the deliberative method they employ: offers use the propose-and-ratify method, while promises use propose-and-challenge. This explains why offers must be explicitly accepted to come into force, while promises only require their addressees to recognize and refrain from challenging them.

Requests are proposals to make joint decisions concerning the *addressee's* actions ("You will  $\phi$ ") using the propose-and-ratify method. The use of propose-and-ratify explains why requests, like offers, must be explicitly accepted to generate an obligation.

Finally, agreements are joint decisions that concern both participants' actions ("We will  $\phi$ "). The bottom row of our table is not split in two because we do not have words in English that distinguish between the propose-and-challenge and propose-and-ratify methods of proposing agreements. But the distinction still applies. Consider two ways of proposing the same agreement: "Hey, do you want to go on a walk?" vs. "Let's go on a walk." The former uses propose-and-ratify, the latter uses propose-and-challenge.

(Our  $3 \times 2$  division suggests one more category: proposals to make joint decisions concerning the addressee's actions via the propose-and-challenge method. My view is that commands and demands occupy this category, but the defense of this hypothesis will have to wait for another day.)

Here's what I find most exciting about this view. The theory of joint decision-making just presented follows from two independently plausible premises: (1) it is possible to make joint decisions that concern only one agent's actions; and (2) we can jointly deliberate in two distinct ways, one that treats proposed joint decisions as accepted by default (propose-and-challenge) and another that treats them as rejected by default (propose-and-ratify). These minimal premises yield a substantive prediction: we should expect to find joint-decision-proposing speech acts that correspond to the boxes in our table above. My conjecture is that we *do* find these speech acts in everyday life—we just don't call them 'proposals in joint decision-making'. We call them 'promises', 'offers', 'requests', and 'agreements'.

If it can be defended, this conjecture would do more than just explain the unity of our four phenomena: it would explain their *diversity* as well. Why do we have these four subtly different normative powers, rather than some bigger or smaller set? The joint decision view provides an elegant answer: because promises, offers, requests, and agreements are the logically possible ways in which we can decide together what to do.

### 3.2. Explaining the Similarities

This paper began by asking: why do offers, requests, and agreements have such similar normative effects to promises? The joint decision view offers a simple answer: because promises, offers, requests, and agreements all result in joint decisions. The normative effects shared by promises, offers, requests, and agreements are the normative effects of a joint decision.

A full defense of the joint decision view would both defend the claim that joint decisions have these normative effects and offer an explanation of why they do. I simply don't have the space to do that here. Instead, I will briefly point out that it is plausible that joint decisions have each of the normative features reviewed in §1.2. (For a more extensive defense of this claim, see de Kenessey 2020: 211–218.) Suppose we jointly decide that I will pick you up from the airport. Intuitively, our joint decision has the following normative effects:

- (1) *Obligation*. Our joint decision places me under an obligation to pick you up.
- (2) *Deontological structure*. Even if I could produce better consequences by volunteering, I am still obligated to pick you up.
- (3) *Directedness*. As a result of our joint decision, I owe it *to you* in particular to pick you up.
- (4) *Waivability*. You have the power to waive my obligation by proposing to retract our joint decision. (More on this shortly.)
- (5) *Exclusionary force*. Once we have jointly decided that I will pick you up, it is inappropriate for me to treat not doing so as a live option.
- (6) *Invalidation by coercion or deception*. If you coerced or deceived me into 'jointly deciding' with you that I will pick you up, this joint decision is invalid and does not obligate me to pick you up.

Supposing that joint decisions do have these normative features, the joint decision view can easily explain why promises, accepted offers, accepted requests, and agreements have them: because they result in joint decisions. Importantly for our purposes, this also explains why promises, accepted offers, accepted requests, and agreements have *the same* pattern of normative features: because they all have the same result, namely a joint decision.



One of the normative features on our list requires further comment: waivability. While it is intuitively plausible that joint decisions can be waived, it may seem that the *way* they are waived is different from promises, accepted offers, and accepted requests. For promises, accepted offers, and accepted requests, the power to waive the obligation is asymmetrically placed in the hands of one party. The promisee can release the promisor from her obligation, but the promisor clearly cannot release herself. However, the power to waive a joint decision appears to be symmetrical. If we jointly decide to meet for coffee, then neither of us has the power to unilaterally waive our joint decision. Instead, we must mutually assent to release ourselves from it. This disanalogy between explanans and explananda presents a *prima facie* problem for the joint decision view.

However, the problem can be solved. Start by noting that retracting a joint decision is no less a deliberative process than making one. To retract a joint decision, one of us will have to propose the retraction, and the other will have to signal their acceptance or rejection of the proposal. We thus face a choice of deliberative method: do we assess a proposal to retract via propose-and-ratify or propose-and-challenge? I suggest that the asymmetry in the power to waive promises, accepted offers, and accepted requests is explained by an asymmetry in the deliberative methods we use to retract their corresponding joint decisions.

Promises, accepted offers, and accepted requests place the parties into two roles: the *agent*, who is charged with carrying out the joint decision, and the *beneficiary*, of whom no action is required. Intuitively, the beneficiary gets the power of release; the agent cannot release herself. I propose that this is because we evaluate the beneficiary's proposals to retract the joint decision using propose-and-challenge, while we evaluate the agent's retraction proposals using propose-and-ratify. If the beneficiary proposes to retract the joint decision, then it is waived by default unless the agent objects.<sup>8</sup> While if the agent proposes to

---

8. Some might balk at the implication that promisors can challenge promisees' release. It may seem that a promisee can release the promisor unilaterally, regardless of whether they want to be released. (Thanks to a referee for bringing up this worry.)

However, it does not seem unintelligible to challenge a promisee's release. Suppose you have promised a talented but overly self-effacing graduate student that you will read her paper. The student, in a moment of self-doubt, attempts to release you from the promise: "You don't have to read it, it's probably no good." You may well object: "I promised to read your paper and you should hold me to it. You have every right to get feedback from faculty." Importantly, if the student then says "Okay, you're right," you remain bound by your original promise.

The objector may point out that the promisee still has the last word over release: if the student insists on releasing you from your promise, you cannot refuse to accept this. But this is compatible with the joint decision view. When two parties reach an intractable disagreement over whether to make or retract a joint decision, they face a choice of deliberative method: which way is the disagreement resolved? As with propose-and-challenge vs -ratify, I suspect we adopt different deliberative methods in different contexts. So, when a promisor and promisee disagree over release, we may adopt a deliberative method that favors the promisee and counts the decision as retracted. This method would give the promisee the last word over release, in keeping with the intuitive data.

retract the agreement, her proposal will only succeed if the beneficiary explicitly accepts it.

There is good reason to deliberate in this way. Often, the point of a joint decision is to strengthen the agent's commitment to a course of action. This committing function would be undermined if it were too easy for the agent to back out. So when the agent proposes to retract a joint decision, it makes sense to use the more conservative propose-and-ratify method, so they cannot back out without the beneficiary's explicit consent. But there is no similar reason to make it difficult for the beneficiary to retract the joint decision. So when the beneficiary proposes to retract, it makes sense to use the more liberal propose-and-challenge method.

In contrast, agreements require both parties to act, and so do not divide the parties into agent and beneficiary. Thus there is no reason to assess proposals to retract agreements asymmetrically. We should evaluate both parties' proposals to retract by the same method, whether that is propose-and-challenge or propose-and-ratify (I suspect both are used in different contexts). So, the reason why it seems that joint decisions are retracted symmetrically is because the most intuitive cases of joint decision-making are agreements. Once we countenance the possibility of joint decisions concerning only one party's actions, we no longer have reason to think that all joint decisions are waived symmetrically. If a joint decision concerns one party's actions but not the other's, then we should expect that asymmetry to be reflected in how it is waived.

### 3.3. Explaining the Differences

While promises, accepted offers, accepted requests, and agreements all have the same pattern of normative effects, there are important differences between them. This section shows how the joint decision view explains these differences. I will pay particular attention to the features that caused difficulty for promissory reductionism in §2.

#### 3.3.1. Promises

Promises are the only member of our quartet that exclusively use the propose-and-challenge method, on which a proposed joint decision is accepted by default unless challenged. This explains two of the most salient ways in which promises stand apart from offers, requests, and agreements: their *apparent unilaterality* and *seriousness*.

First, promises, unlike accepted offers, accepted requests, and agreements, appear to be a unilateral act of the speaker. When I promise to pick you up

from the airport, you don't need to do anything: I simply say, "I'll pick you up, I promise" and I am bound. This makes promises a less intuitive fit for the joint decision view than offers, requests, and agreements, each of which has the call-and-response structure that one would expect to see in joint decision-making (e.g., "Will you pick me up?" "Yes, I will.").

But, as most theorists of promising acknowledge, the promisee plays an ineliminable role in the generation of promissory obligation. I cannot make a binding promise while alone, shouting in the woods. Instead, for my promise to have any force, you must hear and accept it. But crucially, your acceptance of my promise does not require an overt communicative act. If you hear and understand my promise and do not object to it, the promise is binding.<sup>9</sup>

So, promises' appearance of unilaterality arises not because the promisee is uninvolved, but instead because her involvement is tacit and passive: all she has to do is recognize the promise and not object to it. And this is exactly as the joint decision view predicts. For in the propose-and-challenge method, all it takes for an addressee of a proposed joint decision to count as accepting it is for her to recognize it and refrain from objecting.

Let's move to the second feature. While agreeing that accepted offers, accepted requests, and agreements all produce obligations, many readers may still have the lingering thought that the obligations generated by promises are more *serious* than the rest. If I promise to pick you up at the airport (and especially if I say the magic words "I promise"), my obligation may seem stronger, less overridable, than it would be if, for example, I had merely accepted your request to pick you up.<sup>10</sup> This observation seems to cause trouble for the joint decision view. If promises and accepted requests generate obligations by the same means—namely, a joint decision—then shouldn't we expect these obligations to be equal in strength?

The answer to this challenge starts from the observation that some joint decisions are more serious than others, and thus generate stronger obligations to comply with them. Compare a joint decision to meet up for coffee this afternoon with a joint decision to write a book together. The latter is clearly a more serious joint decision, and this strengthens the obligations it generates. It would take less to justify flaking on our coffee date than it would to justify flaking on our book project. There is nothing in the joint decision view that implies that all joint decisions must result in equally strong obligations.

But why would promises *systematically* generate stronger obligations than the rest? The first thing to note is that this is not a general, exceptionless rule.

---

9. See the references in footnote 3.

10. I personally don't feel this intuition strongly: it seems to me that the obligations generated by the four conversations in §1.1 are of equal strength. But I do agree that promises in general have an air of seriousness that isn't shared by offers, requests, and agreements.

For example, my promising to pick you up from the airport generates a weaker obligation than my accepting your request that I be your Alcoholics Anonymous sponsor. Insofar as promises generate stronger obligations than the rest, this is only true *ceteris paribus*.

To see how the joint decision view explains this trend, we must distinguish two functions that joint decisions can serve. (The argument to follow is elaborated further in de Kenessey 2020: 218–221.) Joint decisions serve a *coordinating function* when they help the participating agents to coordinate their actions with each other's actions or preferences. For instance, we may jointly decide to meet at Starbucks at 4pm simply to make sure we end up at the same place at the same time. Joint decisions serve a *committing function* when they help one or both agents to strengthen their commitment to a course of action and communicate that commitment to the other. For instance, a couple's joint decision to be monogamous may serve this function.

*Ceteris paribus*, we should expect joint decisions that serve a committing function to generate stronger obligations than joint decisions that only serve a coordinating function. For the very purpose of a committing joint decision is to produce a strong obligation to comply with it. In contrast, a weaker obligation is often sufficient to coordinate our actions.

I propose that promises tend to generate stronger obligations than accepted offers, accepted requests, and agreements because they are more likely to serve a committing function. When we simply need to coordinate, offers, requests, and agreements are more useful. Agreements most obviously serve the coordinating function, because they place both parties' actions within a single plan. But offers and requests serve another important coordinating function: coordinating one party's actions with the other's preferences. It is more appropriate for me to offer rather than promise when I am unsure whether you *want* me to perform the relevant action—for example, if I'm not sure whether someone else is picking you up from the airport. And so an offer serves to coordinate my actions with your preferences. In reverse, a request serves to coordinate the addressee's actions with the speaker's preferences: by asking me to pick you up from the airport, you inform me that you would like me to do so, and give me a chance to coordinate my actions with this preference. Both offers and requests are well-fitted to serve a coordinating function because of their propose-and-ratify structure: by requiring the addressee to explicitly accept the proposed joint decision, they allow the speaker to propose a joint decision when she is unsure what the addressee wants.

In contrast, promises, because of their propose-and-challenge structure, are not particularly useful for coordination. If all you need is to plan on the assumption that I will  $\phi$ , then often my simply informing you of my intention to  $\phi$  will be sufficient. It makes sense to propose a joint decision by the propose-

and-challenge method when I know that you want me to do something, and I want to assure you that I will do it. And so speakers are more likely to use propose-and-challenge in situations where they need to commit themselves to the addressee. Promises are *ceteris paribus* more serious because they are *ceteris paribus* more likely to serve a committing function. This is how the joint decision view explains the relative seriousness of promises.

### 3.3.2. Offers

The joint decision view explains the difference between promises and offers by appeal to the different deliberative methods they employ. Since promises use the propose-and-challenge method, they generate joint decisions by default unless challenged. In contrast, since offers use propose-and-ratify, they do not generate a joint decision unless and until they are overtly ratified by the addressee.

The feature of offers that the promissory reductionist had trouble explaining was their revocability (§2.1). Intuitively, an offer can be unilaterally revoked if it has yet to be accepted; it is only binding and irrevocable once the addressee has accepted it. The joint decision view explains this observation by appeal to the structure of the propose-and-ratify method. Before it has been accepted, an offer is not a joint decision: it is a proposal to make a joint decision. Offers can be revoked because deliberative proposals can be revoked.

When you and I make a joint decision, we have undertaken a commitment jointly, and so retracting this commitment must be a joint action as well. But when I propose a joint decision to you, that proposal is my individual action, which I have the authority to retract unilaterally. Compare retracting an assertion: after saying, “The birthday party is on Friday” I can say, “Actually, I take it back—I forgot we moved it to Saturday.” Similarly, if you have yet to accept my offer to give you my bookcase, I can say, “Actually, I take it back—somebody just offered to buy that bookcase.” Once you accept my offer, however, it is transformed into a joint decision, and so I no longer have the unilateral authority to retract it.

Promises are not revocable because they typically involve no temporal gap between the proposing of a joint decision and its acceptance. In propose-and-challenge, a proposal turns into a joint decision as soon as the addressee recognizes it and does not object. This leaves no time in which the proposal is the speaker’s to unilaterally retract, for as soon as a joint decision comes into force, the speaker needs the addressee’s assent to waive it. One condition under which a propose-and-challenge proposal could be revoked is if the addressee objects to it. But this seems intuitively right, as illustrated by the following exchange: “I’ll

pick you up at the airport, I promise.” “Thanks, but Susan is going to pick me up.” “Oh, okay, I take it back.”

### 3.3.3. Requests

My main worry about the promissory reductionist’s account of requests was that, by identifying the act of accepting a request with promising, it failed to capture its tight connection with the act of *making* a request (§2.2). Intuitively, accepting a request is essentially a response to a request, rather than an independently intelligible act like promising.

The joint decision view vindicates this intuition. The view says that the acts of making and accepting a request are both moves in the propose-and-ratify procedure of joint deliberation. To make a request is to propose a joint decision regarding the addressee’s actions via propose-and-ratify; to accept a request is to ratify that proposal, bringing a joint decision into force. And so the act of accepting a request is not intelligible independently from the act of making one.

The joint decision view also implies that, though they have similar normative results, accepting a request is not the same act as promising. On the joint decision view, to promise is to propose a joint decision, while to accept a request is to ratify a proposal that has already been made. While both acts have the same result, they are distinct moves in the game of joint decision-making.

### 3.3.4. Agreements

The sticking point for the promissory reductionist’s account of agreements was agreements’ interdependence. If I shirk my side of an agreement, you are no longer obligated to keep yours. The promissory reductionist could only capture interdependence by making our promises mutually conditional, which falsely implies that we can satisfy our agreement-based obligations by both doing nothing.

In contrast, the interdependence of agreements follows naturally if we view them as joint decisions concerning what both parties will do. A joint decision that I will  $\phi$  and you will  $\psi$  is successfully executed only if I  $\phi$  *and* you  $\psi$ . If you fail to  $\psi$ , that is sufficient to make it the case that our joint decision cannot be executed. As a result, our joint decision no longer gives me any reason to  $\phi$ , for my  $\phi$ ing cannot make it the case that our joint decision is carried out. Hence interdependence: one party’s failure to do their part in an agreement makes it the case that the other party cannot bring it about that the agreement is fulfilled, thereby rendering it normatively inert. This explanation does not imply that we



have no reason to fulfill our agreements. So long as it is possible for us to carry out our joint decision, we both have strong reason to do so.

## 4. Objections

I now consider some objections.<sup>11</sup>

### 4.1. Do Accepted Offers and Accepted Requests Really Obligate?

Our first two objections target the opening claim of this paper: that accepted offers, accepted requests, and agreements are similar enough to promises that theories of promising need to capture them. The first of these objections takes aim at the premise that accepted offers and accepted requests generate obligations.

The objection charges that, in at least some cases, the person who makes an offer or accepts a request has the power to change their mind and opt out after the offering/requesting transaction is complete. Suppose that after I offer to pick you up from the airport, and you accept my offer, something comes up that makes it inconvenient for me to pick you up. It may seem that as the offer was mine to make, it is mine to take back: I can call you up and say, “Sorry, something’s come up, I can’t pick you up after all.” Similarly, if I accepted your request to pick you up, I may still have the authority to take my acceptance back.

I don’t fully share this intuition, but supposing it’s correct, it is still compatible with the claim that accepted offers and accepted requests generate obligations. We need to distinguish here between the question of *whether an obligation is generated* and the question of *how it can be waived*. The intuition above concerns how the obligations generated by accepted offers and accepted requests can be waived: in some cases, the objector suggests, the agent has the power to waive their own obligation. But that does not entail that no obligation is generated. To see whether there is an obligation, we need to ask whether, if I had not called you to retract my obligation, I would have wronged you by failing to pick you up. And here the answer seems a clear yes: if I offer to pick you up and you accept, or I accept your request to pick you up, and then I leave you stranded with no further communication, I have wronged you.<sup>12</sup>

What remains to be explained is the intuition that I have the power to waive my own obligation. While this is, if true, a point of contrast with promises, it is one that the joint decision view has the resources to explain. Recall from §3.2 that

11. Thanks to a referee for bringing my attention to the objections in this section.

12. Fruh (2014: 63–66) similarly argues that it is possible to be under an obligation while having the power to waive it.

proposals to retract joint decisions can be assessed using either the propose-and-challenge or propose-and-ratify methods. I argued there that normally, there is good reason to assess the agent's proposals to retract using the propose-and-ratify, so as to make it harder for them to escape their obligation (the agent being the person whose actions are bound by the joint decision). However, this is not an exceptionless rule: there is no in-principle reason why we cannot assess the agent's proposals to retract using propose-and-challenge. And, in a case where we have reason to make the agent's obligation easier to escape—say, in a low-stakes case like giving someone a ride from the airport—it may be appropriate to allow the agent to propose retraction using propose-and-challenge. And that is just what I am doing when I call you and say “Sorry, something's come up, I can't pick you up after all.”

The sense may remain that, in general, the obligations generated by promises are stronger than those generated by accepted offers and accepted requests. But, as I argued in §3.3.1., the joint decision view has the resources to explain this trend.

#### 4.2. Cultural Variability

A second reason for doubting the claim that theorists of promises should be interested in offers, requests, and agreements appeals to cultural variability. The idea is that while accepted offers, accepted requests, and agreements may seem binding in our cultural context, there may well be cultures in which they are not treated as binding, or in which these speech acts do not even exist in a recognizable form. From a cross-cultural perspective, offers, requests, and agreements may seem like heterogeneous categories that are insufficiently robust and unified to require explanation within an aspirationally universal moral theory.

This objection raises an important challenge—but one that applies to promises with equal force. I have not encountered a single paper on promises that adduces any evidence that they are recognized and homogeneous across cultures. In fact, there is a paper arguing that in at least one culture, in the islands of Tonga, people do not use or recognize promises (Korn & Decktor Korn 1983).<sup>13</sup> I see no *a priori* reason to think that offers, requests, and agreements will be more cross-culturally variable or heterogeneous than promises. The only way to tell would be to look at the data. And so to cite cultural variability as a reason to ignore offers, requests, and agreements while continuing to study promises seems to me to employ an unwarranted double-standard. Cross-cultural investi-

---

13. See also Liberto (2016) for arguments that promising is not a homogeneous category. Again, the point is not that these challenges do not apply to offers, requests, and agreements, but that they arise equally for promises.

gation is sorely needed, I agree; but the fact that this need has yet to be fulfilled is not good reason to selectively ignore offers, requests, and agreements.

### 4.3. “Do You Promise?”

A referee points out that the following exchange seems perfectly ordinary:

You: Will you pick me up from the airport?

ME: Sure, I’d be happy to.

You: Do you promise?

The intelligibility of this conversation poses a problem, both for the arguments of §1 and for the joint decision view. In the first two lines, you make a request and I accept it. I’ve contended that this exchange obligates me to pick you up from the airport in a way that is similar to promissory obligation. But if the accepted request obligates, why doesn’t it seem redundant for you to follow up with “Do you promise?” This objection is not limited to requests, either, as it seems similarly intelligible to say “Do you promise?” after an accepted offer or agreement.

However, it also seems intelligible to say “Do you promise?” after some *promises*. Consider:

ME: I’ll pick you up from the airport tomorrow.

You: Do you promise?

This exchange is also unremarkable. Does that mean that my initial speech act did not constitute a promise, or did not obligate me to pick you up? I don’t think so: since one can promise without saying the words “I promise,” I see no reason to deny that my initial statement is a promise. Still, it is intelligible for you to double-check the nature of my commitment by asking “Do you promise?”

This provides a flat-footed response to the objection: as in the case of promising, it is intelligible to follow up on an accepted offer, accepted request, or agreement by double-checking the nature of the commitment that was just undertaken. That is what you are doing when you ask “Do you promise?”

But this response is incomplete. For your question “Do you promise?” does not only serve to confirm the nature of the commitment I already made. If I respond, “Yes, I promise,” this intuitively *strengthens* my obligation to pick you up. How can we explain this, if the accepted request already produced an obligation?

My answer appeals to the claim, defended in §3.3.1, that promises *ceteris paribus* produce stronger obligations than accepted offers, accepted requests,

and agreements, because they more often serve a committing function. On the joint decision view, my accepting your request is enough to make it the case that we have jointly decided that I will pick you up, and thus to obligate me to do so. (This seems right: if our conversation stopped there, I would still be obligated to pick you up.) But joint decisions can be more or less committal, and thus produce stronger or weaker obligations. So the point of your asking “Do you promise?” is to get me to re-propose the joint decision that I will pick you up in a more committal way. Since promises generally propose more committal joint decisions, we mutually understand that by saying “Yes, I promise,” I have strengthened my commitment to picking you up. So, despite initial appearances, the joint decision view can explain the intelligibility of the “Do you promise?” exchange. (For a similar argument, see de Kenessey 2020: 218–21.)

#### 4.4. *Promises to the Self*

I have argued that theories of promising should try to capture the sibling phenomena of offers, requests, and agreements. A parallel argument has recently been made concerning promises to the self. Several philosophers have contended that promises to oneself are importantly similar to promises to others, and so theories of promising should be able to capture both phenomena (see Habib 2009; Rosati 2011; Fruh 2014; and Dannenberg 2015).

One might worry that the joint decision view is a poor fit for this task. For the joint decision view says that promises result in *joint* decisions, and you can’t exactly make a joint decision with yourself.

But this is too quick. The joint decision view has been developed as a theory of interpersonal promises. The question should not be whether this view directly captures promises to the self without modification, but whether it can be plausibly extended to explain them. I believe it can.

The joint decision theorist should say that just as interpersonal promises result in joint decisions, intrapersonal promises result in *individual* decisions. We can then explain the similarities and differences between promises to others and promises to oneself by appeal to the similarities and differences between joint and individual decisions. Advocates of the joint decision view have already emphasized the analogy between joint and individual decisions, arguing that we can better understand some of the normative effects of promises by looking at the parallel effects of individual decisions (Gilbert 2018: 194–97; de Kenessey 2020: 213–14). We might apply this idea in the other direction, exploring how the normative effects of individual decisions can be illuminated by analogy with joint decisions (and thus, on the joint decision view, with promises).

One might object that not all individual decisions are plausibly regarded as promises to oneself. Fair enough: but similarly, not all joint decisions are the result of promises—they might result from offers, requests, or agreements. However, the unified ‘decision view’ I am sketching does suggest that there is no bright dividing line between promises to oneself and other individual decisions. Some might find this objectionable, but I find it plausible. The difference between promises to oneself and other decisions strikes me as a matter of degree. While all individual decisions involve undertaking a commitment to oneself, promises to oneself generate a particularly strong, sticky commitment. Similarly, on the joint decision view, there is no bright dividing line between joint decisions produced by promises and those produced by other means. Instead, there is a *ceteris paribus* difference in degree: since promises are more likely to serve a committing function, they will generally produce stronger, stickier joint decisions. I suspect that we are invoking this feature of interpersonal promises when we use the term “promise” to describe individual decisions that serve the function of making a strong commitment to oneself.<sup>14</sup>

Of course, this is far from a full theory of promises to the self—that would require another paper. But I hope it is enough to show that the joint decision view is not obviously unable to capture this important sibling of interpersonal promises.

## 5. Conclusion

The most important lesson of this paper is that theorists of promising should pay attention to offers, requests, and agreements. What I most hope my readers take away from this paper is curiosity about its question: why do accepted offers, accepted requests, and agreements have such similar normative effects to promises? It matters more to me that my colleagues take up this question than that they adopt my favored answer to it. I have contended that the joint decision view provides an attractive unified account of promises, offers, requests, and agreements. However, the competitors to this account have yet to be articulated. It remains to be seen how other theories of promising—such as the classic trio of conventionalism, the expectation theory, and the normative powers theory—might be extended to capture offers, requests, and agreements. So I’ll conclude by proposing a joint decision: fellow philosophers of promising, let’s get to work!

---

14. Both Rosati (2011) and Dannenberg (2015) emphasize this committing function in their accounts of promises to the self.

## Acknowledgements

This paper grew out of my doctoral dissertation, and so my first and deepest thanks go to my dissertation committee: Stephen Darwall, Julia Markovits, Sally Haslanger, Bradford Skow, and most of all Kieran Setiya, who was the committee's chair. For thorough comments on drafts, I am grateful to David Barnett, Micha Gläser, Tom Hurka, Andrew Lichter, Bernhard Salow, Jonas Vandieken, and two anonymous referees at another journal. For informal discussion, I thank the MIT and University of Toronto philosophy communities. Finally, thanks to the editors and referees for *Ergo*, whose comments led to substantial improvements in the paper.

## References

- Bach, Kent (1995). Terms of Agreement. *Ethics*, 105(3), 604–12.
- Benson, Peter (2019). *Justice in Transactions: A Theory of Contract Law*. Harvard University Press.
- Black, Oliver (2012). *Agreements: A Philosophical and Legal Study*. Cambridge University Press.
- Cupit, Geoffrey (1994). How Requests (and Promises) Create Obligations. *Philosophical Quarterly*, 44(177), 439–55.
- Dannenberg, Jorah (2015). Promising Ourselves, Promising Others. *Journal of Ethics*, 19(2), 159–83.
- Enoch, David (2011). Giving Practical Reasons. *Philosophers' Imprint*, 11(4), 1–22.
- Fried, Charles (1981). *Contract as Promise: A Theory of Contractual Obligation*. Oxford University Press.
- Fruh, Kyle (2014). The Power to Promise Oneself. *Southern Journal of Philosophy*, 52(1), 62–85.
- Fruh, Kyle (2019). Promising's Neglected Siblings: Oaths, Vows, and Promissory Obligation. *Pacific Philosophical Quarterly*, 100(3), 858–80.
- Gilbert, Margaret (1993). Is an Agreement an Exchange of Promises? *Journal of Philosophy*, 90(12), 627–49.
- Gilbert, Margaret (2018). *Rights and Demands: A Foundational Inquiry*. Oxford University Press.
- Gläser, Micha (2019). The Normative Structure of Request. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (Vol. 9, 29–49). Oxford University Press.
- Habib, Allen (2009). Promises to the Self. *Canadian Journal of Philosophy*, 39(4), 537–57.
- Hanser, Matthew (2015). Doing Another's Bidding. In G. Pavlakos (Ed.), *Reasons and Intentions in Law and Practical Agency* (95–120). Cambridge University Press.
- Häyry, Matti and Timo Airaksinen (1998). Hard and Soft Offers as Constraints. *Philosophia*, 18, 385–98.
- Hume, David (1978). *A Treatise of Human Nature*. L. A. Selby-Bigge and P. H. Nidditch (Eds.). Oxford University Press. (Original work published 1739)
- de Kenessey, Brendan (2020). Promises as Proposals in Joint Practical Deliberation. *Noûs*, 54(1), 204–32.



- Korn, Fred and Shulamit Decktor Korn (1983). Where People Don't Promise. *Ethics*, 93(3), 445–50.
- Liberto, Hallie (2016). Promissory Obligation: Against a Unified Account. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics* (Vol. 6, 102–22). Oxford University Press.
- Lichter, Andrew (2021). Promises, Intentions, and Reasons for Action. *Ethics*, 132(1), 218–31.
- Owens, David (2006). A Simple Theory of Promising. *Philosophical Review*, 115(1), 51–77.
- Raz, Joseph (1977). Promises and Obligations. In P. M. S. Hacker and J. Raz (Eds.), *Law, Morality, and Society: Essays in Honour of H. L. A. Hart* (210–28). Clarendon Press.
- Raz, Joseph (1986). *The Morality of Freedom*. Clarendon Press.
- Rocha, James (2011). The Sexual Harassment Coercive Offer. *Journal of Applied Philosophy*, 28(2), 203–16.
- Rosati, Connie (2011). The Importance of Self-Promises. In H. Sheinman (Eds.), *Promises and Agreements* (124–55). Oxford University Press.
- Roth, Abraham (2016). Intention, Expectation, and Promissory Obligation. *Ethics*, 127, 88–115.
- Sánchez Brigido, Rodrigo (2006). Agreements, Promises and Content. *Australian Journal of Legal Philosophy*, 31, 72–94.
- Sheinman, Hanoch (2011). Agreement as Joint Promise. In H. Sheinman (Ed.), *Promises and Agreements* (365–96). Oxford University Press.
- Shiffrin, Seana (2007). The Divergence of Contract and Promise. *Harvard Law Review*, 120, 708–53.
- Thomson, Judith Jarvis (1990). *The Realm of Rights*. Harvard University Press.