

The Normative Force of Promising*

Penultimate Draft — Forthcoming in Oxford Studies in Normative Ethics 6

Jack Woods

1. Introduction

Promises have normative force. The fact that I have promised that I will do something, all things being equal, gives rise to a reason to do it.¹ Why, though, do promises give rise to reasons? I will consider a number of possibilities which don't work (§1.1), then I will sketch a more plausible explanation (§2)—that it is constitutive of the practice of promising that promise-breaking implies liability to blame and that we desire to avoid such liability. This effects a reduction of the normativity of promising to *conventionalism about liability* together with a minimal form of *instrumental normativity*.² I call such a view *quasi-conventionalist*—it makes essential reference to promissory conventions, but the normative upshot doesn't arise from these. It instead arises from the normativity of *desire-based* reasons aimed at properties *grounded* in our conventions.

Quasi-conventionalism can be extended to account for nearly all normativity—excepting the normativity of desire-based reasons—but I will not do so here. My goal in this paper is just to develop a reasonable quasi-conventionalist view of promising. This is independently interesting as well as suggestive of a general reduction of normativity to conventions, desire-based reasons, and minimal instrumental normativity. It is possible to accept my view of promises without accepting my general take on normativity. Likewise, it is possible to accept my view of promising while replacing desire-based reasons with value-based reasons or rational obligations to have certain desires.

* Many thanks to Berit Brogaard, Agnes Callard, Eli Chudnoff, Bradford Cokelet, Catharine Diehl, Erhan Demircioğlu, Cansu Keskin, Benjamin Kiesewetter, Barry Maguire, Daniel Nolan, Mary Robertson, Karl Schafer, Sarah Stroud, Michael Slote, Amie Thomasson, Lucas Thorpe, Lars Vinx, Ken Westphal, and Bill Wringe for helpful comments. Thanks especially to Simon Gaus for useful comments on an early draft and Barry Maguire for forcing me to write my view up. Audiences at the Humboldt University, Boğaziçi University, Middle East Technical University, and especially the Arizona Workshop in Normative Ethics also provided very useful feedback.

¹ I only discuss *reasons*, not *obligations*. The relationship between the two exceeds the scope of this paper. 'All things being equal' here captures the catch-all category of promises which are unhappy because of coercion, contextual cancelation of illocutionary force (such as promising while acting), and the like.

² Readers may notice a more-than-passing similarity to Hart's views on legal normativity as I proceed. One crucial difference between our views is our explanation of reasons to obey a practice when sanctions will not be imposed. Hart uses the notion of an *internal point of view*—a point of view in which we view the rules of a legal system to be binding—to explain why we have reasons to obey the law. I find this mysterious at best; hence my preferred strategy of using *desire-based* reasons and instrumental normativity to explain why we take the rules of a practice to be reason providing. I take instrumental reason and thus instrumental normativity to be understood for the purposes of this short paper. See (Hubin, 2001) for a defense of the innocuousness of instrumental reasons for humane accounts.

Convention-involving views of promissory reasons have three main advantages. First, as Kolodny and Wallace (2003) argue, convention-involving views are the most plausible analyses of promising which avoid explanatory circularity. Second, conventional views are *naturalistic* in that they reduce mysterious-seeming promissory reasons to something non-mysterious such as a conventional practice or expectation. Such views are to be preferred if they can be given. Finally, it seems almost a Moorean fact that promissory conventions and, resultingly, promissory reasons differ from context to context. *Ordnung muss sein*, but not everywhere.³ And some promissory conventions demand that there are few exculpating conditions for promise-breaking. A robust theory of the reasons arising from these domains should accommodate this; doing so is difficult on a non-convention-involving view. I will thus restrict my attention in what follows to views which make use of conventions in explaining promissory reasons.

In addition to satisfying our main desideratum—that acts of promising produce reasons—and the above additional desiderata, a theory of promising ought to explain a number of other features of promissory normativity. First, a theory of promises should explain the *particularity* of promissory reason (§1.2)—that promises give rise to a reason for *me* to keep my promise to *you*. Second, it should be able to explain the diversity of reasons—moral, prudential, political, etc—to keep our promises (§1.3). Third, it should be able to provide a plausible story about cases of conflicting and immoral promises (§1.3). The view I will defend has an easy time with these. Promissory reasons arise from *my* desire to keep my promise to *you*; promissory reasons are distinct from moral, political, and prudential reasons, so we have many different reasons to keep our promises; and, finally, quasi-conventionalism predicts significant (though hopefully outweighed!) promissory reason to keep our evil promises or keep both promises in the case of unexpected conflict but allows that there are, say, moral reasons to break the lesser promise.

My argument in favor of quasi-conventionalism is abductive; quasi-conventionalism can account for cases other conventionalist explanations of promissory reasons cannot, it is naturalistic, it can explain the particularity of promissory reasons, it is flexible enough to explain both types of diversity of reasons to keep our promises, and it can deal with cases of evil and conflicting promises. The putative difficulties with such a view are either overstated or can be assuaged (§2.1-2.2). All in all, it is a tidy and compelling picture of our promissory reasons.

And, importantly, one which significantly better in capturing the sense in which promises directly give rise to reasons. I am not claiming that the reasons which arise directly from the act of promising are always overriding, though I suspect they often are. I am not claiming that we always ought, all things considered, keep our promises, though I suspect often we must. I only

³ “Ordnung muss sein” is true if and only if there must be order.

claim that, all things being equal, when I have promised to do something, then I have a reason, however, small, to do it. This seems a rather central feature of the practice of promising. And promising is a *central* case of normativity. If promises do not provide reasons in all normal non-coerced reasonably well-informed cases, then little is left of the distinctiveness of promising as a central normative notion. Nevertheless, extant conventionalist views of promising fail even this minimal requirement.

1.1 Difficulties for other Conventionalist Views

Consider the following views:⁴

- (1) The practice of promising is a socially useful or valuable practice and breaking a promise contributes to its breakdown. (Cartoon Hume 1972)⁵
- (2) There is an expectation that the promiser will keep their promise and defeating this expectation is a harmful or disvaluable thing. (Scanlon 1990)
- (3) Promises cede normative authority to the recipient of the promise. Breaking promises defies this authority and their interest in our keeping our promises. (Owens 2012)
- (4) We have an obligation to follow those rules the widespread of acceptance of which would maximize good consequences. Promising is such a set of rules. (Hooker 2011)

On each of these views we explain the normative force of promising either by the existence or goodness of a conventional practice (1, 2-4) or by general features of the act of promising which can be viewed as conventional (2).⁶ The primary trouble with (1-4) is that they don't explain why we have reasons to keep our promises in many cases. Consider:

Deathbed Promise: I promise you that I will water your plant after you die. It's an ugly plant; not something the world will be poorer without. Moreover, it's a pain to water the plant. I need to find the water can, fill it, then sprinkle water on

⁴ I am putting aside natural law accounts and 'promises are normative when we should do what we've promised anyways' accounts. These are grotesquely implausible on their face. I am also putting aside Rawls's account of promissory wrongs that trades off of a general moral requirement against free-riding. One reason, mentioned below, is that Rawls's picture has difficulty accounting for the particularity of promissory reasons. Another is that Rawls's picture is famously vulnerable to a circularity complaint (Robins 1984).

⁵ Although this very cartoonish picture of Hume is sufficient for my purposes, there is a more accurate interpretation under which we develop moral reasons to keep our promises (i.e. a settled disposition to disapprove of promise-breaking from a common point of view) as a way of maintaining socially useful promissory practices (Cohon 2006). The resulting view is not miles away from mine, but discussion would be distracting from my present purposes. One notable difference is that, on my account, the reason generator is not a sentiment aimed at *promise-breaking*, but rather a sentiment aimed at a conventionally specified consequence of promise breaking, *blame-liability*. I hope to address the differences between Hume's view and my own in further work.

⁶ Scanlon himself does not take promissory conventions to be essential, but the most plausible development of his view makes essential use of promissory conventions in order to explain the reasonableness of forming an expectation that someone will keep their promises (Kolodny and Wallace 2003).

the plant twice a week. No one knows about my promise. I do not water the plant. It withers and dies. I don't feel guilty or suffer psychological trauma.

There is no plausibility to the idea that violating the deathbed promise contributes to the breakdown of the valuable institution of promising or impinges on my future ability to make promises (1) or that it harms anyone (2).⁷ Clearly utility will be better served by violating some deathbed promises (4). Yet there is still normative force associated with my promising. I can regard myself as having done something I should not have and I would be right to so regard myself; my laziness is not exculpatory. So if (1), (2), or (4) are at best partial explanations of the normative force of promising, (3) is not threatened by our above case, but consider:

No Expectation Promise: I, an untrustworthy friend, ask you for a small amount of money, promising that I will repay you in a week. You have no expectation that I will do so, but you have plenty of money and feel bad for me. You lend me the money. I do not repay you.⁸

In No Expectation Promise, there is again no sense in which you are harmed in the way Scanlon needs—his account of harm trades on the expectation that a promise would be fulfilled (2). Likewise, it's hard to see how such violations could degrade the social institution of promising (1). Just like the fact that we occasionally lie in circumstances where it is expected has no effect on the fact that communication proceeds on the expectation of truthfulness, that we occasionally break our promises has no effect on the fact that promising proceeds on the expectation of typical commitment. With respect to (4), note that utility maximizing promissory practices will plausibly allow us to sometimes break promises when no one depends on their fulfillment.⁹ It would be strange if the optimistic set of rules fails to have an exception clause for this easily rec-

⁷ A reviewer suggests that violating our deathbed promises may inculcate promise-breaking dispositions. In order for this to suffice to get a reason, we need that such dispositions in me will undermine the practice and that I will actually develop such a disposition from even a single act of deathbed promise-breaking. This combination is implausible; even if I would develop a disposition to occasionally break my promises, it's unlikely that my actions have sufficient impact on the behaviors of others that our promissory conventions will break down. Consider, for example, the thought that lying undermines the communicative role of assertion with the fact that we lie often and flagrantly.

⁸ See Scanlon's discussion of the Profligate Pal (1990). His defense against this case relies on my actions violating a norm of gratitude. But although I may violate such a norm in not repaying you, I surely also wrong you by breaking my promise. See Southwood and Friedrich (2009) for further discussion.

⁹ This resembles the famous objection that rule consequentialism inevitably collapses back into act consequentialism (Lyons 1965). Nothing so strong is necessary here though. All we need is that utility maximizing systems of promissory rules will involve many exceptions.

ognizable and obviously harmless type of promise-breaking.¹⁰ Finally, even if you have been ceded normative authority, it's hard to say why I have a reason to repay you; you simply *will not invoke* the authority I have ceded (3). That my action would defy your uninvoked and uninterested authority is not, by itself, indicative of any reason to not perform it. In what sense do you have an interest in my promise being kept?

It might be worried that these promises are only promises in a by-courtesy sense. And I grant that for some promissory conventions such promises don't really count. In other promissory conventions, including the one I implicitly subscribe to, these count as promises and we have reasons, intuitively, to keep them. For those of us subscribing to such conventions, promise breaking in these circumstances is still generally wrong; that there is a reason for me to keep my promise even when there is no expectation that I will and no sanction which will be imposed when I don't.¹¹ As I suggested above, it is important that the general explanation of the normative force of promises is independent of the particulars of this or that promissory convention; it is a mark significantly in favor of a view if it is able to do this, as (1-4) cannot.

In many contexts, such promises do not create expectations of fulfillment, but I would hesitate to say that they carried no normative authority.¹² And, in fatigued enough circumstances, there is roughly no chance that we will invoke our normative authority against, say, politicians, bosses, or parents, even if they have ceded it to us. Perhaps I sometimes have an interest in them keeping their promises—but surely not always. Yet there is a lingering sense in which the politician, boss, or parent had a reason to keep the promise *merely* in virtue of the fact that they had made it. Failure to see this seems to me to be indicative of failing to distinguish between being blameworthy for breaking such promises and having reasons to keep them.

If we accept that we have reasons to keep our promises in both types of case sketched above, then none of the mentioned views will suffice. Or, anyways, views that can capture this intuition are to be preferred. Our abductive argument in favor of quasi-conventionalism can be

¹⁰ Of course, it could perhaps be argued that the optimific set of rules involves rules without exceptions because they are, say, easier to follow. This also seems implausible; it's not as if we have trouble applying our actual rules which have multiple and often very particularistic exceptions. At best, the rule-consequentialist faces an additional explanatory burden here. Thanks to a reviewer for discussion.

¹¹ Note that since my argument is abductive, even if these don't count as legitimate promises—which, of course, they do—there are still very good reasons to prefer a quasi-conventionalist account. Note also that the quasi-conventionalist can not only explain why these count as legitimate promises (for many of us), she can also explain why they don't for others. If they don't count, then there are local conventional conditions on promissory uptake or happiness that exclude these cases from counting as situations in which we can successfully promise.

¹² David Owens's early work accounts for these sorts of cases in terms of *bare wrongdoing*, explained in terms of “non-instrumental value in its making sense for your deliberations to be shaped by promises and other obligations.” (Owens 2006) This strikes me as less promising than the quasi-conventionalist story I tell; for one thing, it is hard to explain in a compelling way how such value can be non-instrumental.

strengthened by noting another feature of promising that (1) and (4) neglect—the *particularity* of promissory reasons.

1.2 Particularity of Promissory Reasons

There is something personal about someone breaking a promise to us. It is not just that they have offended against the community of promise makers and keepers; it is not even that they have offended against everyone or the world by making it a bit worse off. They have offended *us*. Kolodny and Wallace (2003) put the point well:

... the practice view cannot account for our conviction that the wrongness of breaking promises is to be explained by reference to the reasons that the promisee *in particular* would have to object to the promiser's failure to perform. (emphasis mine)

Everyone has reasons to object to someone breaking a promise *when* the wrongness of promise breaking is understood in Rawls' or Hooker's terms.¹³ But this is not how we understand the wrong of promise breaking. We think that there are special reasons to keep *our* promises and that claims against *me* as a promise-breaker are local; they are not reasons for everyone to ensure my promise is kept or claims that others can make.¹⁴ A theory that neglected this feature of our reasons to keep our promises would be unacceptably revisionary. This is not to say that there aren't sometimes reasons for *me* to make *you* keep your promise. Often there will be. And it is not to say that others never have any claim against me for breaking my promise to you. Often they do. But this isn't sufficient to assuage our intuition that I have a special reason to keep *my* promise to *you* and that when I don't, *you* have a special claim against *me*.

Quasi-conventionalist accounts have little trouble explaining why the claim against the promise-breaker is local to the promisee; (social) rules of the practice set the conditions under which and, more importantly, *who* may impose or prescind from imposing the sanction. Typically, for most promises, this is the promisee. In some cases, others have a right to sanction as well, but this right is often lesser and if the promisee lets the offender off the hook, others are blocked from sanctioning. And quasi-conventionalism directly captures the sense in which *I* have

¹³ Their remark is, of course, directed at Scanlon. Nevertheless, the point applies broadly. Hooker might object that it maximizes utility to have objections be limited to only the victim of a promise-breaking, but why would this be true? Wouldn't it be more useful overall to have the entire society able to criticize a promise-breaker?

¹⁴ One special case here is the witnessed vow—think of marriage vows. In such cases, a far larger group of people share the claim against a promise breaker. I set oaths, vows, and the like aside for simplicity here.

a special reason to keep *my* promise to *you*. I care about *my* culpability. In particular, I care about *my* culpability to *you*,¹⁵ since typically you are the one who has a claim about me I care about.¹⁶

Accounts like (1) and (4) tend to miss this point; the reason they provide is a reason for everyone to get me keep my promise as well.¹⁷ If this affects me more than others, it's because I'm typically in the best position to make myself keep my promise. There *are* such general reasons, of course, but they do not suffice to explain the particular reason I have to keep my promise, which generate merely by promising. This simply is not a reason available to others.

So, the correct view will imply that legitimate promises, in any normal case, give rise to reasons to keep them and, moreover, that these reasons are local to us in the sense that they are *our* reasons, not reasons available for everyone or, even, everyone who is part of the promissory convention. Again, this is not to deny that there are cases when promises or apparent promises do not give rise to reasons.¹⁸ But these cases should be relatively rare and have to do with conditions for a legitimate promise to not come off or, in Austin's nice phrase, be *unhappy*.

1.3 Extrinsic Reasons to Keep Promises and Further Benefits

We also want to allow that there are other *extrinsic* reasons to keep our promises; reasons over and above the special reasons which arise from the mere fact that we have promised. In my view, much of the confusion that has arisen about promising comes from the fact that there are often *many* different reasons to keep our promises. Some are extrinsic—the expectation that we will perform the promise (Scanlon) or more general moral concerns (Hooker and others)—and some are intrinsic—such as promissory reasons. As the cases above show, extrinsic reasons can be re-

¹⁵ See below for more discussion of the distinction between desires to avoid being blame-culpable and desires to avoid being blame-culpable to you *in particular*. It is plausible that in the majority of cases, both are present.

¹⁶ Simple modifications of the view are possible which do even better at capturing this feature. We could complicate the picture by introducing a relative notion of blame-liability—relativized to the permitted sanctioners—and claim that our reason to not break our promises stems from our desire to not be blame-labile to them. A fully developed version of this view needs to accommodate the point that our desires are often a desire to not be blameworthy *to you* for *doing this or that*. This means that I will typically have more reason to keep my promise to my mother or my best friend than I do to a random person. And I will perhaps have more reason to keep my promise to co-sign a loan than I do to meet you for lunch—'perhaps' since extrinsic reasons also matter here. But this strikes me as a virtue of my account—I simply do have more reason to keep my promises to those I tremble at being blame-labile to. I still need a desire that holds for any case. But it seems plausible to me that we typically have a *generic* desire to not be blame-labile to anyone for anything, and thereby a *generic* reason to keep our promise to *anyone* for doing *anything*. This is plausible—see the discussion below—but fully justifying it will require further work. Thanks to Agnes Callard for pushing me on this point.

¹⁷ Kolodny and Wallace *op cit* show that Scanlon's account (2) likewise has trouble meeting this constraint.

¹⁸ Cases of coercion, cases of failures of information, and suchlike may cancel our reasons to keep our promises. Or, perhaps, such "promises" do not count as promises at all. Nothing turns on which we say.

moved without removing the intrinsic reasons to keep our promises.¹⁹ The temptation to place all the normative force on the extrinsic reasons is strong, but it gives an inadequate picture of the normative contours of promising.²⁰ One tempting diagnosis of the mistake common to all these views is that they conflate extrinsic reasons for certain promissory conventions and consequences thereof with intrinsic reasons arising, at least partially, from the convention itself. This is like arguing that *en passant* is illegitimate in chess because chess without *en passant* would be a more elegant game. Focusing on extrinsic reasons yields another benefit. Consider:

Conflicting Promises I promise to pick you up from the airport and I also promise to meet my friend for coffee to discuss some tragedy in their life. I have no antecedent reason to believe that there will be any problem keeping both. Your plane gets delayed a tremendously long time due to an unpronounceable volcano exploding, making it impossible for me to fulfill both promises. I had no reason to expect that this would happen. You, savvy traveler, will have relatively few problems finding your way to your hotel, but my other friend is quite distraught. I cannot ask your permission to break the promise; there is not enough time. I text you that I will not be picking you up and comfort my friend.

In this case, I want to say that (a) I have a reason to keep both promises—in fact, I think I am *promissorily obliged*, at least in the sense of the formal ought attached to our promissory conventions, to keep both promises, (b) I have an extrinsic *moral* reason to break and keep the promises I did, and (c) given that I had to break at least one promise and I broke the obviously less important promise, it wouldn't be cricket to blame me, though it would be within your rights to do so. So we have a nice case of conflict between *moral* reasons and *promissory* reasons. Assimilating the sort of reasons generated by our promises to moral reasons results in trouble capturing the lingering feeling that something is less than ideal about my actions. This is better captured by keeping my reasons to keep my promises separated from my reasons to do what is right.

Quasi-conventionalism promises an elegant solution to this troubling problem. How can it be that I have done the best thing I could (morally), yet nevertheless have done something wrong? It would be wrong, morally, to blame me for acting this way if I am not in any significant

¹⁹ Liberto (this volume) is sensitive to this point, but she posits a flat ambiguity in the normative force of promising—sometimes promises give rise to expectation-based-obligations, sometimes they give rise to normative-authority-based obligations, etc. But in addition to the positing of an unnecessary ambiguity, this view runs into a rather difficult problem. How should we know which obligation we incur by making a promise? Liberto suggests that we can individuate by focusing on what counts as breaking the promise, but this just shifts the bump in the rug. If what counts as a promise breaking of the promise made is sourced in the speech act performed by the promisee, we need an ambiguity account of a seemingly unified type of speech act. One wonders why we should call both types of speech act promises at all. If the act is unified, but the obligation varied, then it is a mystery how the act could give rise to the obligation all by itself, as it seems to. Liberto's view, in spite of its ecumenical virtues, is too costly.

²⁰ There is also the typical temptation here to confuse *reasons to have a practice* with *reasons internal to the practice*. There are many moral and prudential reasons to propagate a promissory convention, but these do not yield a reason for me to keep my particular promise since, as has been noted here and many times elsewhere, breaking my particular promise does not undermine the socially valuable promissory practice. See Maguire and Woods (Unpublished).

sense culpable for making inconsistent promises. The answer is that I have done something for which I can be blamed, by promissory lights, and hence have done something *wrong*, by promissory lights. But I am immune to *moral* blame. If our promissory reasons are weighty, then we get even closer to explaining my sense that there is something seriously wrong with my breaking a promise, even when I am morally blameless for so doing, as in the above case.²¹ An exactly similar move can be made with respect to immoral promises—when I promise to kill my victim with your favorite knife, I go against my promissory reasons by killing them with a corkscrew, though I should morally have not killed them at all.²²

Summing up, we have found a number of desiderata for our account of promising.

- (A) It should account for how promising nearly always gives rise to reasons.
- (B) It should explain the *particularity* of promissory reasons.
- (C) It should explain, in a natural way, how we can have moral, prudential, political, etc. reasons *as well as* promissory reasons to keep our promises.

The best approach to satisfying (A-C) identifies a feature of the conventional practice of promising which (a) places the promise-maker in a unique position, (b) gives a unified picture of promissory reasons without undermining the extrinsic reasons to keep our promises, and (c) holds even in the cases of promising sketched above. We can then use this feature to explain why we have reasons to keep our promises, even though the feature I will identify as doing this work doesn't *by itself* generate reasons to perform a promise. The feature I have in mind is *blame-liability*.

2. The Positive Proposal — Blame-liability, Blameworthiness, and Blame

Consider the Drones, a famous, yet fictional, gentleman's club in London. The Drones have a number of by-laws, including the rule that one must attempt to steal a policeman's helmet whenever the opportunity presents itself. Failure to do so has the sanction of having the entire club throw dinner rolls at the offender. Chester Fotheringly-Bantam, a veteran member, sees a clear

²¹ Of course, we would still have to give an account of how promissory reasons conflict and of how we can be wrong in breaking a promise even when we could not but do so. The theorist who accepts moral dilemmas might not find these arguments compelling. Well and good, but moral dilemmas are *puzzling!* My account of the promissory case of a moral dilemma captures the important sense attached to them while dissolving the puzzle.

²² An even more radical move is to suggest that the picture I have sketched for the reasons arising from promises holds *mutatis mutandis* for reasons arising from morality. If morality is largely conventional, then we need to ask the question of why we ought to be moral. The best such reason I can think of is that we generally care about being moral or, at least, to not violate moral rules. If this is right, then it is to be expected that our promissory conventions can conflict with our moral conventions, which is not to say that there aren't moral rules about keeping promises. There often are—think Kantian-flavored positions—but there sometimes are not—compare act-consequentialist pictures. One thing the consequentialist picture seems to get intuitively wrong is that there's a non-negotiable push of some type to keep our promises. One thing the Kantian picture gets wrong is that there is often a non-negotiable moral push to break them. Separating these two practices splits these problems nicely—and this recipe can be applied to other puzzling cases of normative conflict.

attempt to steal a policeman's helmet, worries about the consequences because of his fledgling career as a barrister, and consequently fails to attempt to steal it. Chester is pelted with dinner rolls at the next club dinner. Here we have a sanction justly imposed.

But now consider Boffy Slantingwood, a neophyte Drone. An opportunity to steal a policeman's helmet presents itself; Boffy, naif that he is, fails to recognize the opportunity; and then a discussion ensues prior to the next dinner about whether to pelt Boffy with dinner rolls. It is decided that Boffy is new and inexperienced enough that his violation of the by-law isn't really deserving of the sanction. The hardline members of the Drones insist to a man that Boffy is liable for the offense and they shouldn't make special exception for new members, especially since they pelted poor Chester just last week. Chester, of course, agrees with this. The more relaxed members of the Drones agree that Boffy is liable and concede that if they were to pelt Boffy with dinner rolls, no one could really complain. Yet they point out that (a) the threat of being so pelted is probably sufficient and (b) they're short on dinner rolls. At the next dinner, the members rise to pelt Boffy, but at the last minute they drop their bready missiles and sternly warn Boffy that he will not be so lucky next time.

In the second case, what we have is a case of *sanction-liability*, but without the imposition of the sanction or even sanction-worthiness. Boffy has violated the rules of the Drones club and it is thus within their purview to pelt him with dinner rolls. But note that the fact that it is in their purview to do so does not by itself imply that they must or that it would be reasonable to do so. They have discretion about the imposition of the sanction. There were good reasons to not pelt him, but, if they were to pelt him, he would have little recourse. He could argue that they shouldn't since he was new; but this isn't to say that they weren't within their rights to do so, just that doing so was unkind or unnecessarily by the book. We could go into some detail about why they had good reasons to not impose the sanction in this case—lack of other dinner rolls, that Boffy wasn't clearly not an irreverent buffoon (the quality that the helmet rule is meant to cultivate), etc. But it is sufficient for our purposes to note the distinction between being open to sanction, being sanction worthy, and the sanction being imposed. Between, that is, being-pelted-with-dinner-rolls-liable, being-pelted-with-dinner-rolls-worthy, and being pelted.

For another nice example of *sanction-liability*, consider the law. Someone can be open to legal sanction but imposition of the sanction would be unreasonable. It is by no means obvious that good legal systems will always give judges no discretion in imposition of sanction. There is thus a distinction between being legal-sanction-liable and being actually sanctioned, even in the case where the legal system is functioning correctly. There is also a corresponding difference between it being *reasonable* or *good* to impose a sanction and it being *open* to you *qua judge* to do so. If we imagine the purpose of the law is to maintain order and reduce harm, and that recognition of

this is part of participating in the practice of executing the law, then imposing sanctions for crimes where the impact on the citizenry is low, when doing so is likely to provoke civil unrest, is *bad* and *unreasonable*, even though it is in the purview of the law to do so.

The case of the law is instructive in another fashion. There are many ways for a legal sanction to be good or reasonable. One clear way, as just mentioned, is for the sanction to fit with the purpose of the law—if the purpose of the sanction is to deter others from breaking the law—as in the specious justification of the death penalty—then it is reasonable to impose the sanction when it deters. And permissible even when it does not, at least as far as the law goes.

Sometimes the reasons are extrinsic to the practice. When keeping a promise would result in a moral tragedy, there are tremendous *moral* reasons to avoid keeping the promise. It is thus (morally) unreasonable to blame someone for breaking the promise in that circumstance. And sometimes it is a balance of the two. Even if the purpose of a prescriptive grammatical rule is to prevent confusion, then there still may be a small reason to cite the mistake in contexts where it will cause no confusion—we should not develop bad habits. But posh supermarkets which have ‘10 items or fewer’ signs are simply obnoxious as the social reasons to speak with the vulgar are far more important than a constant grammatical peccadillo.

So, ultimately, we have a threefold distinction between the imposition of sanctions, it being good or reasonable to impose sanctions—subdivided between intrinsic aims and extrinsic aims—and being open to sanction. Certain conventional practices will collapse the first and the third. In certain legal systems, for certain crimes, if you are open to legal sanction, the court will have little to no discretion in whether to punish you.²³ But not always. What counts as good or reasonable with respect to imposing sanctions may be local to the convention, but it often is not. Reasonableness arises here from other domains—morality, for example—and from the purpose of the practice in the first place. But *liability* for legal and promissory sanction is completely local, being entirely set by the rules and laws of the practices in question.

I take promissory conventions, like club rules and legal conventions, to be constituted by sets of rules and corresponding sanctions which can be imposed. Perhaps some conventions are not so constituted, but the cases under consideration certainly seem to be. A detailed discussion of the nature of conventions would take us afield, so let me set out briefly what I take to characterize the sort of conventions at issue. They are a trio of rules, corresponding sanctions, and (sometimes) an internal purpose that generates a notion of reasonableness in imposing sanctions, as in the above Drones case. Some of these rules set the condition of who may impose or lift the sanction when appropriate (rules of adjudication), as is most important for the case of

²³ Three strike laws are a good example. Judges often actually have the most discretion with respect to the amount of sanction, but sometimes their hands are tied even here—compare mandatory minimums in American drug law.

promising, along with conditions for introducing new by-laws (rules of change) and deciding whether a formal rule has been broken, as in the case of the Drones. These rules tend to be internalized into our practice of recognizing certain individuals as having this authority—and they have this authority because we accord it to them and accept that they have it. They are, in the Hart’s sense, social rules.

Rules may be explicit or implicit, but the conditions for breaking them must be roughly clear enough to allow a participant to the convention to judge, generally, whether or not the rule has been violated. The rules may be laid down in explicit form with all possible cases considered or defined elsewhere (‘dress is white tie’), or they may have a more vague and particularistic feel (‘dress is business casual’). All that is really necessary is that rules be such that either there is a determinate property of liability for sanction that can be determined by participants to the convention or there is a rule governing who may decide whether someone is sanction-liable.²⁴ Note, –liable, not –worthy.

Blame-liability is exactly the sort of property we need to explain the normativity of promising. It is present in every case of promise-breaking, it is intrinsic to the promising convention itself, and it is plausibly something we have reason to want to avoid. For example, cases like deathbed promise and no expectation promise are outliers where the ordinary extrinsic reasons do not obtain, and in which the promise-breaker is not clearly blameworthy for doing so, yet we feel compelled to say that there is still a reason to keep our promise. We thus need a feature that is present in both cases in our explanation of our reasons in this case. But blame-liability *is* present in both cases. Even though I *will not be* blamed in either case by the promisee (because they do not care or are not alive), and am perhaps not even blameworthy (watering is a pain and no one cares about the money), I am still liable for blame.²⁵ And, if liability to blame for promise breaking is sufficient to generate a reason for me to keep my promise, then we have reason to keep our promises even in the outlier cases I sketched above. But is liability for blame enough?

²⁴ For a nice recent discussion of conventions which attempts to accommodate Lewisian conventions as well, see (Marmour 2009). The tradition of taking promissory conventions to be constituted by constitutive rules dates to Rawls’s early discussion of rules (Rawls 1955). Note that there does not seem to be a rule of recognition, in Hart’s sense, for promissory conventions. Well and good, unlike the law, it is unclear that we need to have a validating rule such as the rule of recognition on top of rules of change and adjudication. This topic gets more messy (and more interesting) in the context of full-bore quasi-conventionalism; I hope thus to address it elsewhere.

²⁵ This is not to say that in ordinary such cases, I will never be blamed by others. Bystander members of the promissory convention, at least in many cases, can sometimes blame (morally, prudentially, and sometimes promissorially) for breaking my promise to someone.

2.1 From Liability for Blame to Reasons to Keep our Promises

The key to the plausibility of the quasi-conventionalist story is a claim about typical human desires. It is that, typically, rational creatures care about being liable for blame.²⁶ Note that this is not the claim that we do not like getting blamed; it is rather the claim that we do not like being *liable* for blame. Put aside for now the question of whether or not this is true. If it is true, then if our desires give rise to *desire-based* reasons and if we have instrumental reasons to carry out actions that realize our desires, we have a *pro tanto* reason to not break our promises. What we need for this linking claim to go through is that if we want it not to be the case that A and doing X will bring about A , then we have a *pro tanto* reason to avoid X -ing. This seems to be a very plausible minimal claim about the desire-based reasons.²⁷

If one accepts this connection between desires and reasons, given the claim that we care about being blame-liable, and given that in all typical cases of promises, breaking them implies the breaker is liable for blame, we have reasons to keep our promises in such cases. If I have made a legitimate promise, then I will be blame-liable if I break it. I do not want to be blame-liable, so I have a *pro tanto* reason to avoid breaking my promise. How strong a *pro tanto* reason? This will depend on the strength of my desire to avoid being blame-liable. Ideally, it should be a strong desire, but there is no guarantee of this. But, I think, this is not much of a problem. So long as we typically have a strong desire to avoid being blame-liable, we will get a correspondingly strong reason to avoid breaking our promise. I do not, however, claim here that our promissory reasons to keep our promises are strong, even though I think it's true. My purpose is to construct scaffolding, not yet to strut it.

The desires I am focusing on here are generic—they are desires to not be blame-liable [to anyone in particular for anything in particular], not a desire to not be blame-liable [to Mark FOR sending my paper late] or to not be blame-liable [to humanity FOR buying a selfie-stick]. That such generic desires exist is antecedently plausible—consider the analogous case of desires for food. The other sort of particularized reasons exist, of course. We may want to eat, but not want to eat the meal placed in front of us. Likewise, I think, we might want to not be blame-liable, but, in fact, not care a whit about being blame-liable to Joe for not showing up for lunch.

²⁶ I mean the notion of desire here to be fairly broad notion—something like an end or a psychological state with a world-to-mind direction of fit.

²⁷ Note two things. First, there are resources for downplaying worries about such humean accounts (Schroeder 2007), including a way to defend the idea that some of our reasons swim under the veil of unmentionability since they are too minimal to be mentioned in the context of explicit practical reasoning. Second, it is possible to swap out a non-humean account of desire-based reasons and run the same argument—this especially since my desire to not be blame-liable is nothing like a desire to chug a bucket of paint.

Perhaps Joe broke his promise to me last week, so I don't feel particularly worried about him blaming me.²⁸ Is this a problem with my account?

No; in fact, it is a benefit of it. I, presumably, have less reason on balance to keep my promise to Joe than I do to my saintly mother who has never violated a commitment to me. My account captures this nicely since my *on balance* reasons to keep my promise to Joe are less weighty. Note, though, that this does not show that I am not blame-liable to Joe for breaking my promise; I am. I am simply not, perhaps, *blameworthy* for breaking this promise and my reasons to keep it, though existent, are weaker than they would otherwise be. The interaction of particular and generic desires thus allows a more nuanced treatment of our practical judgments.²⁹

That generic and particular reasons interact in this way can be made additionally plausible by noting that when we don't seem to care about being blame-liable to another human being—which I think is rarer than it might seem—this has often to do with cases where we would not be blameworthy for breaking our promise to them because they have wronged us in some way. Given that Joe has just broken his promise to me, it would be unjust for him to blame me for breaking mine to him; I am, perhaps, not blameworthy for breaking this promise (given other details easily spelled out.) I might then not care about being blame-liable to him for breaking mine; I have an easy and conventionally acceptable defense of my actions *to him*—two wrongs do, sometimes, make an okay, at least to the wrongdoer—but I have no defense to someone else for breaking my promise to Joe—I ought to be better than him. The most plausible story about this seems to me to be that I have a generic desire to avoid being a promise-breaker (to avoid being blame-liable), no particular desire to avoid being blame-liable to him for breaking my promise, but a desire, perhaps, to avoid being blame-liable to another for breaking my promise to Joe.

We have identified a feature of promise-breaking that is present in all cases of happy promising; this means that our worrisome cases provide no counterexample to us having reasons to keep our promises so long as we care to avoid blame-liability. Put slightly differently, this

²⁸ Note that much of what I say could be maintained *even if* we do not possess these generic desires, as long as, for arbitrary people and actions, we typically possess the particularized desires. I think this is also very plausible—thinking otherwise seems to me to, once again, conflate the difference between being blameworthy and blame-liable for violating a commitment—but I won't argue for it here since I think the posit of generic desires makes for a more plausible and interesting account.

²⁹ Note that we might worry that the generic desire is *composed* of particular desires (a worry like this is suggested by a reviewer), but consider again hunger and the interaction of my strong desire for food and repulsion at the particular food in front of me. These are seemingly distinct desires interacting with each other. Note also the contents of the generic desire are aimed at possessing a certain property—being blame-liable—whereas the contents of the particular desires are aimed at ways in which we might possess this property—being blame-liable to *a* for X-ing. It seems fairly clear that we can and frequently do, perhaps with a bit of irrationality, want to be some way, but not want to be any of the ways of being that way; even more clear is that we can and frequently do want to be some way without wanting to be that way in some particular way. In such cases, I hold that the desires weigh against each other; the one does not eliminate or obscure the other. So, in particular, they both produce reasons.

means that there is no situation in which we, if we are humanly-typical, do not have reasons to keep our promises, unlike the alternative conventionalist views mooted above. But it also means that if we are not humanly-typical—if we do not care about being blame-liable—we have no reason. It is to this problem that we now turn.

2.2 Worries for Quasi-Conventionalism about Promising

The obvious problem with this account is that it makes our reasons to keep our promises contingent on our desires. What do we say about the person who simply does not care about being blame-liable? One simple response to this, first articulated by Philippa Foot, is simply to say that such people are odd.³⁰ *We* do care about being blame-liable. People who don't care about such things are odd and we should not worry overly much about their presence. For them, perhaps there is no reason to keep a promise, outside of extrinsic reasons and fear of sanction, but what of it? Most of us are not like that and the existence of a few such people or even their possibility will not make useful social conventions like that of promising implode.³¹

But something more can actually be said, I think, about the fact that we typically care. Consider someone who is unsporting. The soccer player who takes a dive at any opportunity or the bridge player who sneaks a peek at the opponents's cards when they're not looking. Such things do not invalidate the game, though they generate liability to sanction.³² Suppose such players care about the sanction, but not the liability for sanction. Intuitively, there is something wrong with them. They've done something wrong. But why? The most plausible explanation is that we think that there's something wrong with breaking the rules or implicit norms in these contexts, *even if we're not caught, even if we're not actually in danger of sanction*. Of course, where we think that there's something wrong with doing something, we typically desire to not do it.³³ We don't want to be the unsporting card player or the dive-taking soccer player.³⁴

³⁰ See her unfortunately recanted-upon (Foot 1972). I develop this line of response elsewhere.

³¹ In fact, I view the fact that such people do not have promissory reasons to keep their promises (they may have prudential reasons, of course!) as a feature, not a bug, for reasons much like those laid out in (Manne 2014). A detailed discussion would take us afield here, but I address the issue in other work (Woods Manuscript).

³² Sometimes these are actually not rule-governed in the sense of being explicitly forbidden by the rules. This is a reason to go with blame-liability instead of simply the property of not breaking the rules, More on this below.

³³ This desire is potentially entangled with the view that someone who does such a thing, including ourselves, is to be badly regarded. This causes no problem for my point; it is a potential explanation of the grounds of our desire to avoid blame-liability.

³⁴ Note that this concerns the typical interactions of a pair of psychological states, nothing more. I make no claim (here!) about the rationality or irrationality of thinking something wrong yet not wanting to do it.

That's one further thing that can be said. Here's another. Consider the difference between a voluntary promise and a voluntary response to a threat—say, deciding to pay a blackmailer. When I promise, I voluntarily undertake to do something and do so on threat of sanction. When I cave to threats, I undertake to do something on threat of something bad happening to me. Of course, the promissory sanction is regulated by a set of rules and the threat typically not, but we can fix this up easily enough. Consider a ruling body that has imposed their will on a people by force and imposed an extraordinarily rigid set of “sanctions” on their governed. Intuitively, I have more of a reason to obey my promise than I do to obey the dictates of such a governing body. Why? Presumably, it is because I care about being open to sanction for promise breaking—I view such a sanction as legitimate.³⁵ Whereas I do not care about being open to sanction from the governing body—just whether I'll be sanctioned and the likelihood thereof.³⁶

For a final reason, consider the distinction between two sorts of reasons to obey laws. In Boston, it is against the law to jaywalk. Nevertheless, Bostonians jaywalk flagrantly. The only possible reason to not do so is the almost unimaginable threat of actually getting a ticket. When, as very occasionally happens, someone gets ticketed, we view it as a great injustice, though one that was within the (legal) rights of the cop to create. This is a case where there is no legal reason to not jaywalk except for how much we care about the sanction and the likelihood of the sanction being imposed. Likewise where the sanction is likely to be imposed, but where we view the sanction as unjust; think of marijuana laws or public drinking laws.³⁷ But in other cases, where there is almost no threat of sanction, we nevertheless think that we have a reason to obey the law. Voting laws are somewhat like this; as are tax laws. Even though sanction is rarely if ever imposed, we nevertheless think that we have a reason to obey these laws. Even when we think that the benefit of us actually obeying the law is roughly nil—think of voting in a relentlessly blue state. Why? Again, I want to suggest that we care about being in violation of these laws.

³⁵ Again, note that this is a claim about the interactions of certain psychological states, nothing more. The reader might incline to the view that it is sufficient to generate a reason that we recognize a certain practice as legitimate; this is very similar, again, to Hart's views about the normativity of the law. But I find this entirely mysterious. Why is it sufficient that we recognize that someone has legitimate authority over us for us to have a reason to obey such authority? My explanation is clear—because we typically also *care* about not trampling legitimate authority or authority we have internalized. The alternative view seems to me to either need to accept this connection as brute—undesirable—or explain why we have such reasons—unpromising. I favor my style of explanation.

³⁶ Clearly this is related to the fact that I have *voluntarily* undertaken the threat of sanction in the promise case, but not in the governing body case. But this is not sufficient as I may have very good reasons to voluntarily undertake subjugation by a despotic government, yet the point still holds.

³⁷ You might worry here that the important part is us not being blameworthy here. But we might be blameworthy to jaywalk—jaywalking is tremendously dangerous, so we have extrinsic prudential reason to not jaywalk and can be blamed for doing so.

This is admittedly very impressionistic, but I think the point is a solid one. There is simply a difference between recognizing a practice and its sanctions as legitimate and caring only about the sanction that might be imposed. When we recognize a practice and its sanctioning practice as legitimate, we care to not violate it and we go to great lengths to defend ourselves from our failures to comply—think about someone who has not voted in an election trying to justify themselves against someone who claims they have a duty to do so. When we accept and internalize a set of norms, we take following these norms seriously.³⁸

You might think that there really are two distinct properties here. One is the property of having violated a rule or implicit norm. The other is the property of being liable for sanction based on your violation of a rule. We might care about one, but not the other. But I think these run close together, especially in the case of something like promising where the sanction is something like blame and resentment for violation of the rules. The content of the blame is that we broke the rules, in a way that is not generally acceptable. When the sanction is tied so closely to the fact that you have violated the rule, then there really is very little gap between caring about being sanction-liable and being in violation of the rule.

And in cases where they come apart, we may care or not about violating the rule—yet we nearly always care about being blame-liable. To care for obeying the rules in the absence of caring about being blame-liable would be incredibly rule-*fetishistic*—in the sense rightly disparaged in (Smith, 1994). We are not typically like this. But to care about being blame-liable without caring about the rules strikes me as more reasonable. It is more or less the attitude sensible folks take towards ordinary etiquette.³⁹ Note that if I had to abandon the notion of blame-liability in favor of, or even add in, caring about simple rule-breaking as a source of reasons (we can be fetishistic, after all), my account would not totter—as long as there is, as there obviously is, a distinction

³⁸ Maybe the reader will worry that we often have reason to change our norms. Perhaps, but that is not relevant. Etiquette norms have pull long after we recognize how arbitrary and, quite frankly, silly many of their dictates are. Part of the problem is that if we do not uniformly change our practice, then we might be accepting a more stringent set of requirements than we are inclined to enforce. And this seems unfair. So there is sometimes also extrinsic reasons to stick with a non-ideal set of norms. Note that these are reasons for our practice to be the way it is, not practice-internal reasons arising from what we care about. Even if there are reasons for us to change our practice, that has little relevance to the reasons arising from our desires to be good practitioners. All it means is that these reasons might be outweighed. But it *by no means* implies that we have no reason to obey the rules of our current promissory practice. Just as the fact that it would be better to eradicate norms of etiquette which demand excessive politeness outweighs, but does not undermine, our etiquette-internal reasons to be overly polite.

³⁹ Note that a striking example of this is when someone is not in a position to blame us through rule-regulated behavior—perhaps they have lost their status as a blame-leveler due to some sin. Typically in such cases, we lose much of our reason to avoid breaking the rules as we lose blame-liability. See also the discussion above of cases where we lose *blameworthiness* but not *blame-liability*. The particulars will depend on the particular contours of our promissory conventions.

between breaking a rule and being blameworthy for breaking the rule. But I think blame-liability better captures what we typically care about.

Finally, note that we need to distinguish between caring about being blame-liable and caring about being blameworthy. It is *prima facie* plausible that we care even more about being blameworthy than we do about being blame-liable. But this is insufficient to explain our full desires with respect to conformity to a practice and immunity to sanction. We prefer to not be blameworthy, but we would also rather be entirely clean of the possibility of blame. And this is a very natural state. Recent work by (Keupp, Behne, and Rakoczy 2013) shows that young children (ages 3 and 5) absorb normative practice so thoroughly that, even in explicitly goal-directed tasks, they protest actions that do not perfectly mimic the method they had been shown of achieving the goal. That is, they *over-imitate* in their own behavior and protest non-imitation.

I have assumed our desires give rise to reason regardless of their reasonableness.⁴⁰ If this is right, then we should expect that as we get better at identifying instrumental reasoning, such as that involved in goal-directed behaviors, we cease over-imitating, but maintain a concern with acting correctly in tasks which are more method-directed, such as moral, promissory, etiquette, and game-oriented practices. But we, even on reflection, know quite well that in many such cases, we will not be blameworthy in virtue of not contravening any purpose such practices have...yet we regard ourselves as having reason to act correctly. This corresponds nicely with the findings of the study just cited, when the focus was placed on the reactions of older children.⁴¹

Importantly, this helps to further explain why we care about *sanction-liability*, not only *sanction-worthiness*. Blameworthiness, as I have sketched it above, depends directly on intrinsic or extrinsic aims of a practice. We are not *worthy* of blame merely by contravening a practice. Something further needs to be said to justify leveling a reactive attitude like blame at us. But many domains of normative concern, though the rules governing them have an extrinsic goal, do not have an *obvious* built-in intrinsic goal in the way that shooting a pool ball or, even the law, has a built-in goal. It thus very plausible that when it comes to such domains, it is unnatural to care only about be-

⁴⁰ Note again that I could still tell the story I want if I assumed that our desires needed independent justification in order to generate reasons as, very plausibly, promissory practices are a very valuable part of our social behavior. I prefer not to do this, but my account does not rule it out.

⁴¹ Of course, as above, many of these desires might be directed at people explicitly breaking the rules; perhaps our desires to conform to certain practices is very particularized. It's difficult to extract the sort of distinction between generic and particular desires given this sort of experimental data, especially given the authority role of the experimenter. However, my point here is merely to suggest why we nearly always have some desires to avoid being blame-liable, even perhaps to avoid being blame-liable *a for X-ing*. The justification of generic desires would require additional argument, such as that given above, for the usefulness and antecedent plausibility of having desires to avoid having the general property of being blame-liable to anyone for anything. Again, note that we can plausibly invoke the distinction between blame-liability and blameworthiness to explain difficult cases. Often the reasons we don't seem to care about being blame-liable to someone for something has to do with the unreasonableness of *them* blaming me legitimately.

havior that serves the point of the norms governing it. Rather we care about behavior that is in conformity with the rules for the practice, encapsulated in liability to sanction.⁴²

3. Conclusion

To summarize the above discussion: If we care not only about being blamed for breaking our promises or even being blameworthy, but also about being liable for such blame and if our desires give rise to reasons, then we have reason to keep our promises, even in outlier cases like those I started with. That we *do* care about being blame-liable, while contingent, is in many ways to be expected. Besides the psychological results I cite, being blame-liable is often tightly connected with violating the rules of a practice that you regard as legitimate; when we regard such practices as legitimate, we care not only about being punished for stepping outside of the rules, but also about keeping within the rules themselves. The practice of promising is exactly a case where we typically regard the practice as legitimate, care not only about being blamed, but also about being open to being blamed, and this is because we do not want to put ourselves in a position where we cannot defend our actions against the accusation that we broke the rules (which is in itself a type of sanction when we care about breaking them.)

If this is on the right track, then there is serious plausibility to our style of explanation. We can explain the special reasons that attach to promising by exploiting our desires and the fact that practices governed by social rules, such as promising, set conditions for being liable for sanctions such as blame. Such reasons have the features we want: they arise from just about all cases of promising, they make essential reference to *our* promise breaking (they are reasons *for us* to avoid breaking our promises to *you*), they allow extrinsic reasons to keep our promises (resulting in a more adequate picture of our reasons to keep our promises as well as dissolving puzzles about evil and conflicting promises), the explanation is flexible enough to allow reasonable contextual variation in promissory reasons, and they are entirely unmysterious. Quasi-conventionalism about promising, I conclude, has quite a bit going for it. More than enough for us to believe it.

⁴² Barry Maguire and his perceptive class have convinced me that I could recast much of what I say in terms of *thin blameworthiness* incurred simply by breaking a rule and *thick blameworthiness* incurred by breaking a rule that matters. As little turns on how we cast the point, I will continue with my terminology.

Works Cited

- Cohon, R. (2006). "Hume on Promises and the Peculiar Act of the Mind." *Journal of the History of Philosophy*, 44(1): 25-45.
- Foot, Philippa (1972). "Morality as a System of Hypothetical Imperatives." *The Philosophical Review*, 81(3): 305-316.
- Hart, H.L.A. (1961). *The Concept of Law*, Oxford: Clarendon Press.
- Hooker, Brad (2011). "Promises and Rule-Consequentialism." in *Promises and Agreements: Philosophical Essays*, Oxford University Press.
- Hubin, D. C. (2001). "The Groundless Normativity of Instrumental Rationality." *The Journal of Philosophy*, 98(9): 445-468.
- Hume, David (1978). *A Treatise of Human Nature* (LA Selby-Bigge, Ed.). Oxford University Press
- Keupp, Stefanie, T. Behne, and H. Rakoczy (2013). "Why do children overimitate? Normativity is crucial." *Journal of Experimental Child Psychology*, 116(2): 392-406.
- Kolodny, Niko, and R. Wallace (2003). "Promises and Practices Revisited." *Philosophy and Public Affairs* 31(2): 119–154.
- Liberto, H. (forthcoming). "Promissory Obligation: Against a Unified Account." in *Oxford Studies in Normative Ethics*
- Lyons, David (1965). *Forms and Limits of Utilitarianism*. Oxford University Press.
- Maguire, B. and Woods, J (Manuscript). "Explaining Epistemic Normativity."
- Manne, K. (2014). "Internalism about Reasons: Sad but True?" *Philosophical Studies*, 167(1): 89–117.
- Marmour, Andrei (2009). *Social Conventions*. Princeton University Press.
- Owens, David (2006). "A Simple Theory of Promising." *Philosophical Review*, 115(1): 51-77.
- _____ (2012). *Shaping the Normative Landscape*, Oxford University Press.
- Rawls, John (1955). "Two Concepts of Rules." *Philosophical Review* 64(1): 3-32.
- _____ (1971). *A Theory of Justice* Cambridge: Harvard University.
- Robins, M (1984). *Promising, Intending and Moral Autonomy*, Cambridge University Press.

Scanlon, Thomas (1990). "Promises and Practices." *Philosophy and Public Affairs*, (3): 199-226.

Smith, Michael (1994). *The Moral Problem*. Wiley-Blackwell.

Southwood, Nicholas, and D. Friedrich (2009). "Promises beyond assurance." *Philosophical Studies* 144(2): 261-280.

Woods, Jack. "Footing the Cost (of Normative Subjectivism)" (Manuscript)