

The Authority of Formality*

Jack Woods

⟨⟨ Standard “this is a draft, treat it as such” caveat here.⟩⟩ This is an attempt to work through a number of ideas and the suggestion at the end needs to be tightened up and made more pointed. However, in the aim of getting things out there for useful feedback, I put more material than necessary in the draft for people to poke at. The main account can be found in the early half of §3 and the important use of some recent psychological results to explain away the felt categoricity of putatively substantive obligations is at the close of that section. At least skim §2 before reading these; it will be difficult to see why I attack this problem the way I do without this discussion. (apologies to those, like David Copp, Evan Tiffany, and Steve Finlay, whose pictures should be discussed more extensively here. There will be a section distinguishing the particulars of my view from theirs in the final draft.)

1 Introduction

Metaethicists often distinguish between *substantive* obligations and *merely formal* obligations (see Copp 2004, McPherson 2011, Parfit 2011, and references therein). The latter are supposed to be mere standards of correctness, without a necessary connection to our reasons, whereas the former are supposed to have a definite connections to what we really truly ought to do. Parfit glosses the distinction in terms of two senses of normativity: the *rule-implying* sense and the *reason-implying* sense. Etiquette is normative in the rule-implying sense—it lays down criteria for correct action *by the lights of etiquette* whereas morality is normative in both senses—it lays down criteria for correct action *by the lights of morality* and, crucially, correct moral action implies that we have (some) reason (simpliciter) to do it.

The distinction could also be, and sometimes is, glossed in terms of guidance or authority: substantive obligations are intrinsically *authoritative*, in that they are supposed to guide the actions of all rational agents (Tiffany (2007, pg. 248) defines authority *for an agent* in exactly this way.) That we morally ought not steal the backpack of the distracted Syrian refugee is supposed to guide our actions, period, though human frailty being what it is, it may not.¹ Merely formal obligations, in contrast, only guide action when supplemented by something like the desirability or value of satisfying the obligation. For example, that we ought to chew with our mouth shut is supposed to guide our actions only insofar as we care (or should) about politeness. Simply by itself, our politeness-based obligation to not expose our mouthful of half-mashed banana is not supposed to guide our actions.²

*Thanks to Derek Baker and Barry Maguire for useful discussions of this material and to the audience at the Uppsala Language and Metaphysics of Normativity conference for useful feedback on a handout-sized draft of this material.

¹I do not mean here that we should deliberate on our moral obligations in deciding how to act—that would, of course, run afoul of the paradox of problems similar to the paradox of hedonism (see Railton (1984) for the classic pertinent discussion.) A looser, but necessary connection between our moral obligations and guidance would suffice for the distinction being made here.

²One needs to be sure here to prescind from instrumental justifications for caring about politeness. “by itself” means by itself, not “in combination with typical human affect” or what have you.

We could go on listing ways theorists have attempted to get at this distinction, but the preceding should suffice to illustrate the general idea behind this common distinction. And it is a very common distinction, playing a role in our theorizing about the normative and informing our actual deliberative practice. We take substantive obligations seriously in a way we do not take merely formal obligations; we criticize others more easily when they violate their substantive obligations; we *care* about substantive obligations more than we do the merely formal ones.

Although this distinction is very common, there is no accepted explanation of it. That is, there is no accepted explanation why some types of obligation, like morality, are authoritative, but others, like etiquette, are not. Given how obvious this distinction has seemed to many, this is rather problematic. Tristram McPherson, discussing Scanlon's domain-relative conception of reasons, points out that

Almost all of us take reasons to be normative in a way that contrasts with other formally normative systems such as the rules of chess or the *schmeasons* standard...
...metanormative realists need to explain why reasons [points in favor of doing something favored by the correct norms formally governing action] are *robustly* normative, whereas *schmeasons* [points in favor of doing something favored by the inverse norms which also formally governing action] are merely formally normative (McPherson 2010, pg. 233)

and the demand holds generally, regardless of whether we think in systems of norms giving rise to reasons or, as we have so far preferred, in terms of obligations. I think this demand hold even for those of us who deny that there are any really categorical types of obligation—we need to explain *why* we so frequently have thought there are.

My aim is thus to reconceive *what it is* for a (type of) obligation to be authoritative, drawing on a general theory of normativity I have begun developing elsewhere (Woods 2016, Maguire and Woods ms.) I will develop a view on which no types of obligation are substantive in the above sense, but where there is nevertheless an interesting difference between these types of obligations. In yet another formulation, I do not deny that we might always have reason to do what we morally ought, but I do deny that our reasons to do what we morally ought follow from the putative categorical nature—in the interesting sense of that term—of our moral obligations. In my view, there are only merely formal obligations backed by weighty, robust, and constant reasons. This is a reconception, not a rejection, as I will demonstrate that there are systematic differences between putatively substantive and putatively merely formal types of obligation; the difference, however, is not that one is basically reason-implying while that the other type is not; it is rather a difference in how we view these types of obligation. Put pithily, the substantive obligations, the ones which have authority, are the ones we take to have authority.

Call an obligation *subscripted* when it arises from a system of norms that we describe with an adjective like 'moral', 'epistemic', 'etiquette' and so on. I reject that any subscripted obligation to ϕ implies, in the absence of side conditions, that I have reason to ϕ . I will not argue for this claim here. The aim of this paper is, rather, to tell an alternative story that captures (at least most of) what we wanted from substantive obligations in the first place. Since this is the aim, I start by characterizing what the supposed *intuitive* difference between substantive and merely formal systems of norms is; what we expect a substantive obligation to be like. I suggest that none of the obvious candidates for distinguishing these systems of norms will work to cleanly divide intuitively merely formal obligations from intuitively substantive obligations. The best

we can do is a hybrid intuitive characterization.

Turning then to merely formal oughts, I lay out a number of important and often overlooked features of merely formal systems of norms. I use this discussion and the prior discussion to lay out a list of desiderata for an account of the difference between merely formal and substantive obligations—they should be similar enough, in the sense of having the same basic explanation, but they should be different enough that we can explain why we distinguish the two cases and why we take substantive obligations to be as they are. The account I lay out below satisfies these desiderata, fits with some interesting findings considering the psychology of norm-internalization, and fits with a generally attractive research program I have been developing in other work. The key move in closing the gap between my picture of substantive obligation and what we take substantive obligations to be like draws on features of the psychology of norm-internalization which have been studied by Keupp, Behne, and Rakoczy (2013, 2016). Call a norm *task-oriented* when there is an explicitly articulated or easily seen goal the norm aims at achieving. For example, instructions for how to tie our shoes are task-oriented—they are aimed at resulting in tied shoes. Instructions for how to treat one another morally are not—the aim of morality, insofar as there is one, is a matter of substantive sociological discovery and not immediately accessible to us ordinary folks. Keupp, Behne, and Rakoczy’s discovery, about the differences in how we view *task-oriented* and non-*task-oriented* norms, furnishes a plausible explanation of why we have taken substantive obligations to be special. Roughly, we take non-*task-oriented* norms to be the sort of thing we should follow to the letter; as we do not see that there are tasks associated with them, we do not allow that their purpose can be satisfied by any number of means. These facts provide a compelling explanation of our view of substantive obligations and, thus, underwrite the idea that we get what we really need out of the reductive story I tell below.

Before getting into the nitty gritty here, let me recap my argumentative strategy. I will not argue that the existence of truly substantive *subscripted* obligations, like morality, is incoherent or problematic. What I will do instead is show that an account that denies that there are any substantive systems of norms and thereby grounded subscripted obligations (a) can say something useful and interesting about this difference and, (b), that what we say about this explains why we *would* view certain obligations as substantive, even though none of them really are. I have argued elsewhere that this background picture is attractive (Woods 2016, Maguire and Woods ms). The work here strengthens the case in favor of this background picture by showing that certain obstacles to it, such as the fact that it implies that neither morality nor epistemic obligations are substantive—is less of an obstacle than one might have thought.

1.1 Reasons and Obligations

One of the confusing issues in approaching this issue has to do with the differences in each theorists’ particular approach, as is clear from the quote from McPherson above. In discussing Scanlon’s view, he is pointing to the very difference we are attempting to elucidate, but he puts it in terms of reasons instead of obligations. Others, including Scanlon, have talked about the differences between normative *domains* or normative *outlooks* or something similarly vague. No one can be faulted for this; it is a fundamental issue and such issues nearly demand such vagaries. To make the below discussion as precise as possible, I will spend some time here regimenting the discourse. When I talk about a system of norms, I will mean a set of forbiddings and permissings, as well as favorings and disfavorings, of actions, and with perhaps some locally defined internal concepts local to the system of norms. I presume that a system of norms will give rise to a *subscripted ought*: the property of actions that are maximally in accords with the

favorings, disfavorings, forbiddings and permissions of the system of norms. I will sometimes refer to an action ϕ that is part of a set of actions maximally in accords with a system of norms d as d -fitting. In particular, an action is d -unfitting when it is no part of a set of actions maximally in accords with a system of norms d .

I will use the adjective ‘substantive’ to characterize both systems of norms and the obligations they induce. Strictly speaking, this should always be interpreted in the former usage a system of norms which gives rise to a substantive obligation. This should cause no confusion. I will not carefully distinguish between ‘ought’ and ‘obligation’. Strictly speaking, the latter word is more appropriate when we are talking about an obligation which is grounded by a system of norms, whereas ‘ought’ works for both this case and the ought simpliciter of most reason. Again, this should cause no confusion.

The rules of tic-tac-toe embody a system of norms. One is permitted to mark any open square on one’s turn. One is forbidden from erasing or taking back one’s move. One is forbidden from making two moves at once. Player X may only mark an ‘X’, Player O may only mark an ‘O’. etc... A *mark* is an ‘X’ or an ‘O’; each player may only have one mark. A *win* is a configuration of the board where there are three consecutive ‘X’s or ‘O’s. The winner is the player whose mark is on the board in three consecutive squares. The loser is player who is not the winner. A *draw* is a configuration of the board which is not a win. Etc...The aim of play is to configure the board into a win; a move is favored insofar as it advances this goal. Etc. Even this trivial example is only a partial spelling out of the permissions, forbiddings, favorings and disfavorings of a game so notoriously simple as tic-tac-toe. Reality is much more complicated. But it is enough to say that we have reason_{tic-tac-toe} open with a corner play and that we ought_{tic-tac-toe} block any two-consecutive pattern unless we can win immediately by playing elsewhere. To do otherwise is *incorrect*_{tic-tac-toe}; playing in an incorrect fashion is *tic-tac-toe-unfitting*.

This way of understanding systems of norms sits nicely with the contemporary orthodox semantics for *ought* where, roughly, background information, such as a system of norms_X, gives rise to a ranking of the possible ways things could go (possible worlds) in terms of how well they accord with X.³ I ought_X to do something just in case the X-norms induces a ranking on worlds such that those in which I do it rank above worlds in which I do not. We need not presume that this ranking is *complete* in the sense that any two worlds are comparable, nor that it is *maximal* in the sense that there is a best-ranked world. Further details would distract and make commitments irrelevant to my overall point. It *is* worth noticing here that this semantic story is fully compatible with every system of norms being substantive, none being substantive, and any mediate position. This semantic story only provides an account of what is is for a system of norms to provide a standard of correctness.

My aim is to characterize then reconceptualize the sense in which certain systems of norms are taken to be substantive. Importantly, both my characterization and my reconceptualization will treat a system of norms being substantive in terms of it implying *reasons*, in the most generic sense of considerations in favor, to act as that system of norms directs. I will not argue here that these generic or flat normative reasons, which I will sometimes call *unsubscripted* reasons, are themselves substantive. The reason for this is that claims like “I can see that, all things considered, I have most reason to do this, but why should I do it?” and cognate constructions

³See (Kratzer 1977) for origins of this account and (Chrisman 2015, ch. 2-5) for a useful contemporary overview from a philosophical standpoint. We sidestep worries arising from differences in the various *flavors* of ‘ought’ since, in the cases under consideration, they are all more or less agentive ‘oughts’, with different ordering sources.

strike me as conceptual confusions.⁴ Insofar as they make sense at all, they seem to me disguised claims that we have no reason to do anything, that reasons don't exist, or the like. Insofar as there is anything like a reason—and our ordinary modes of practical deliberation are sunk if there are not—there will also be a generic notion of a practical, contributory, normative reason. What these reasons are is, of course, a matter of substantive philosophical dispute. As is what it is to act in a way which is based on this reason. I will abstract from these worries here, intending a fairly broadly ecumenical picture of what these reasons are and how we can act in such a way that our actions are based on them.⁵

It is a further and important question whether there actually exists an *unsubscripted ought* or an ought simpliciter.⁶ This would be an ought that is not itself relativized to some background system of norms or standards which can interpret our flat considered judgments about what we ought to do and arbitrates cases of conflicts between subscripted 'ought's.⁷ Given our unsubscripted notion of reason or favoring, we can define an unsubscripted ought by the recipe that we ought to do what we have most (unsubscripted) reason to do. We should note at the outset that the semantic story given above is consistent with there being such a thing and, more importantly, for many different theoretical recipes for how to generate such a thing. The semantic story for the unsubscripted 'ought' just takes the grounds of our unsubscripted notion of a reason to fill the *X*-role and set up the obvious system of standards accordingly.

It is the easiest of tasks to construct some notion of ought simpliciter given these materials. If we took value to explain or be fundamental in generating our unsubscripted reasons, then we ought simpliciter to do those things which were value-maximizing (or value-increasing, value-promoting, etc.)⁸ If we start with a deontological picture where the findings of reason were, in some sense, connected with decisive reasons, then we interpret the ought simpliciter in terms of rationality. We could then treat reasons simpliciter as contributory reasons to promote rationally-fitting behavior, or, alternatively, we could say that there are only decisive reasons (things that bring us in accord with the findings of rationality). If we were reasons primitivists (god forbid), then we would get the unsubscripted notion of 'ought' immediately out of the connection between ought and most reason without any need to invoke the grounds of the reason. Finally, if we take desire to explain our unsubscripted notion of a reason, then we have reason simpliciter to do those things which advance our desires and we ought to do those things which are maximally desire satisfying. Similarly, even if there are multiple grounds of equal weight for our reasons, say epistemic, prudence, and moral reasons, which did not always weigh nicely, then we can construct a *partial* ought simpliciter in terms of what these reasons all agree upon.

For my purposes here, I don't care which underlying explanatory or metaphysical story we

⁴In (Woods 2014) I argued that expressions of moral judgments could coherently be conjoined with a denial of any evaluative attitude taken towards them; I was careful to distinguish this case from all-things-considered ought judgments. Toppinen (2015) emphasized this, pointing out that when we make incredibly salient that it's the all things considered ought judgement being made, it really does sound confused. I hope to explore the connection of my earlier work to this claim about judgments of reasons simpliciter elsewhere.

⁵See Sinclair (2016) for a useful expansive account of how to act on a reason which avoids the putative "elusive reasons" problem.

⁶See (Baker this volume) for worries about the existence of this kind of 'ought'. In a way, the Humean position I am going to construct below concedes his worries. I think, however, it does so in an illuminating way which does not "change the subject".

⁷Note I say 'can', not 'does'. I mean to allow that we might mean something else in ordinary judgments which display an explicitly unrelativized 'ought'. The important thing is that there is a notion of an unsubscripted 'ought' or notion of an unsubscripted 'reason' which could interpret how we break ties, describe what we all things considered should do, and the like.

⁸Using maximization here is a simplification which does not affect the overall point; all I commit to here is that we can treat reasons and obligation here in terms of each other and that we can make sense of both within the sort of Kratzerian semantic framework sketched above.

tell about the notion of reason simpliciter which, in turn, explains the ought simpliciter. What I care about is that there is a recognizable notion of ‘reason’ and ‘ought’ which, if it existed, would serve to stand as “true” normativity (in the sense of Broome (2013)) and to serve to ‘close deliberation’ (in the sense of Schroeder (2011)). This is not to claim that there are such oughts or reasons, of course, but just to claim that we recognize what they would be if there were such and that such a notion of reason structures our ordinary practical deliberations.⁹

My task here is not to defend the existence of a generic or flat notion of a normative reason (see McPherson (this volume) for steps in this direction), but a few remarks might help to see where I’m coming from. Our practical deliberations typically proceed by means of subscripted reasons. We talk about what prudence favors, what morality favors, what rationality favors, and the like. It seems, though, that no matter which norms we invoke, we can always ask whether we should do what they favor. In other words, constructions like:

(OQR) Morality, Prudence, Rationality,... all favor doing it, but should I do it?

always seem reasonable questions, no matter which system or systems of norms we plug in for ‘...’. If that’s right, if OQR is a schema that is always coherent, then there has to be a notion of a generic normative reason that we are appealing to in order to make sense of OQR in its various instances. I take the generic notion of a reason to be, in a sense, the limit of OQR, and I take the question of what it is to be a substantive philosophical question.¹⁰

This claim about reason simpliciter does not imply that there is always an answer to questions like OQR. It is only a claim that the question makes sense, even if it is unanswerable. One potential response to my presumption is to claim that the ways of making sense of it differ from context to context and from sets of subscripted norm to sets of subscripted norms. This, however, itself demands an explanation and I’d suggest that whatever explanation can be given of this fact could itself be massaged into a notion of an unsubscripted reason. For example, if someone of Humean persuasion were to suggest that it is the fact that we care more about this or that system of norms on this or that occasion, then this itself strongly suggests that we explicate the notion of reason simpliciter to our desires (in fact, this is more or less the explanation I favor, though I don’t presume it below.) Likewise, if someone suggested that the overall value of breaking it this or that way in this or that context were what we were asking after in OQR and that the salient values could shift from context to context, then we could explicate the notion of a reason simpliciter in terms of promotion of value, allowing that values shifted from context to context. And so on.

Another reason, an important one, for favoring the existence of reason simpliciter has to do with the obvious commensurability of certain normative conflicts. If nearly all of my subscripted obligations tell me to ϕ , but my mild affection for elegance favors avoiding ϕ -ing, it seems clear that I should ϕ . This is easily explained with the existence of reason simpliciter; it seems nearly impossible to explain otherwise, especially as we complicate the case to get conflicts between sys-

⁹I disagree here with Baker (this volume), who claims that the notion of an ought simpliciter is simply to vague or metaphorical to do the work people have wanted. There may be no such reason, but it seems manifestly clear we presume as much in our typical practical deliberations.

¹⁰Derek Baker, in personal correspondence, has suggested that there is a bespoke context-local notion of a generic reason which plays the role in each individual case of something like OQR. Though this dodges my thought, I think it does so at the significant theoretical cost of making our practical deliberations unsystematic. Steven Finlay, in his 2014, offers a theory of the interpretation of all things considered ought-judgments that fits nicely with the sort of picture I myself favor. There are differences, but an elongated discussion of them would distract here.

tems of norms. However, it is important that reason simpliciter should not arbitrate all conflicts between systems of norms. Sometimes our reasons really are just on a par (Chang 2012). Below, I will use this as a desiderata for our account of the difference between substantive and merely formal systems of norms, but it is also a desiderata for a reasonable view of reason simpliciter.

Turning back to systems of norms, we make no presumptions about the distribution of strict and loose features of a system of norms. A system of norms may consist entirely of a list of forbiddings and permissions with no non-trivial favoring and disfavorings.¹¹ It may consist entirely of non-trivial favorings and disfavorings where we concoct obligations in terms of what we have most reason to do and to avoid. It may consist almost entirely of an *aim* which implicitly gives rise to both favorings and disfavorings—things which would further impede the aim—and forbiddings and permissions—things which would definitely frustrate the aim and things which do nothing of the kind. Finally, it might be entirely particularistic, having no explicit rules or aims, but where we participants have a sense of which things are favored and which things disfavored (or, forbidden and permitted). Any of these systems of norms can be massaged into a ranking and will give rise to a notion of ought in the above sense.

With all this in mind, we will talk about the difference between substantive and merely formal obligations in relatively neutral terms. I will talk as follows: *subscripted obligations* arise from *systems of norms*. These obligations will either be *substantive* or *merely formal*, insofar as as the claim that we *ought_X* to do something implies, in the absence of side conditions, that we have *reason simpliciter* to do it.

I hold that there are no substantive obligations in this sense. One way for this to be true is that there is no such thing as a reason simpliciter, so we never have reason to do what we *ought_X* because we never have reason to do anything period. This is not what I am claiming. What I claim is that even if there were reasons simpliciter, we would not have reason simpliciter to ϕ just because we *ought_X* to ϕ .

2 Obligations, Substantive and Formal

To develop my account, I need to discuss the intuitive conception of substantive and formal obligations. I have two main aims. First, I will explore what we have in mind when we claim that a system of norms and the obligations it induces are substantive. I will aim at giving an intuitive characterization of this difference, not a philosophical one, though it will connect up with the definition of substantiveness I just offered (in denying that any subscripted obligations were substantive). I will then turn to laying out some contours of merely formal systems of norms and obligations, articulating some of the crucial features of merely formal norms which will play a role in our reconceptualization.

2.1 Substantive Obligations

Since there is widespread agreement that there is a difference between substantive and merely formal obligations, we ought to be able to find a distinguishing mark. Perhaps the most salient of the putatively authoritative domains of obligation is morality, so we should expect to be able to find some perceptible difference between morality and other systems of norms that seem to

¹¹‘Non-trivial’ since we can always cook up favorings and disfavorings given a system of forbiddings and permissions by claiming that we have decisive reason to do anything we’re forbidden from not doing, decisive reason to avoid doing anything we’re forbidden, and so on.

obviously be merely formal. Our task in this section is thus to try to make sense of what could justify our claim that morality is authoritative while systems of norms like etiquette are not. In the service of that, I will start by canvassing a few immediate suggestions for marking the difference between formal and substantive obligations that do not work, then I will attempt to articulate what I take the best suggestion for doing so to be.

2.1.1 Directiveness

Let a system of norms that articulates not only constraints on appropriate action, but also particular suggestions for how to behave, *directive*. Morality, at least as sometimes construed, is *directive*. Epistemic norms are also plausibly directive—we are told not only that our beliefs ought to be apportioned to the evidence, but that we ought to actively apportion our beliefs to the evidence. One suggestion then for what distinguishes substantive obligations from merely formal obligations takes DIRECTIVENESS to be characteristic of substantive obligation.

This is also a natural thought since some theorists have taken substantive obligations to be *guiding* in the sense that they should inform our practical reasoning about what we ought to do. Finlay, for example, writes:

We can agree with the contemporary consensus that experiencing the consideration that p as normative, or as a reason to ϕ , is to experience it as ‘counting in favour’ of ϕ -ing. But the vital feature here is that the experience of normativity is essentially an experience of autonomous authority. ‘Autonomy’ can mean a lot of different things; here I mean only that normative authority is not alien to the thinking self or self-determining agent. Experiencing something as normative for you forestalls sincere declarations of indifference or skepticism about its practical relevance, of challenging ‘So what? What does that matter?’ (Finlay 2007, pg. 229)¹²

However this notion of guidance is to be cashed out, it cannot be in terms of directedness. Paradigmatic merely formal obligations, such as etiquette, are also directive, favoring treating people with respect. Paradigmatic substantive obligations, such as rationality (construed in terms of standards), are plausibly not directive as they merely lay down constraints for acceptable belief.

Of course, views about the contours of these systems of norms may differ, but the point is that there are recognizable systems of norms which would be of the substantive variety which are not directive and there are recognizable systems of norms of the merely formal variety which are directive. So this route to distinguishing substantive from merely formal obligations is a non-starter. There is something to be said for *guidance* as a criteria, but we will return to this below.

2.1.2 BUT I DON’T CARE ABOUT THAT **and** IN FORCE

We can also reject, nearly immediately, one attempt at the difference between merely formal and substantive obligations. Kant, famously, seemed to suggest in his Groundwork that we can distinguish between instrumental and categorical obligations in terms of their surface grammar (or, rather, in terms of the *form* of the ought-judgment). The suggestion here is that since we comfortably say “you ought to do your homework *if you want to pass*”, but “you ought to avoid hurting others *if you want to be good*” feels uncomfortable, the former is instrumental, whereas

¹²See also Tiffany (2007), Korsgard (1997), and Joyce (2001). Finlay writes slightly later of this view that it is obvious and I agree, though the details of how it is to be guiding is a difficult matter.

the latter is categorical. The problem, as pointed out above, is that demands of etiquette are categorical in form, but they are a paradigm case of the merely formal (Foot 1972). Even if there were this difference in form, of course, it would be a significant jump to the claim that the grammar of the form of various ought-judgments revealed something deep about their status as authoritative. This jump has not, to my knowledge, been made in a way which would underwrite the widespread agreement about the substantiveness of the moral. So the GRAMMATICALITY mark wouldn't, by itself, do the work we want even if it were actually a mark. Which it ain't.

A natural refinement of Kant's thought which avoids Foot's worry treats the categoricity of morality in terms of our response to claims about obligations. It is acceptable to respond to Emily Post's claim that you ought to bring a gift to the party by declaring that you simply don't care about being rude. Responding to the airing of a moral obligation by claiming to not give a rat's ass about morality doesn't seem similarly acceptable. Of course, this perception does not clearly correspond to deep difference between these systems of norms. Suppose we were playing pool in a bar and someone pointed out that the shot I just made was illicit by WPA rules. Since the rules in force are bar rules, I'm totally entitled to tell them to take a walk, that I don't care about WPA rules (in this context). Likewise, if they point out that the shot I made was illicit by bar pool rules, then it seems rather weird, at best, to say that I simply don't care. Bar pool rules are in force.

There are also non-game oriented formal systems of obligation that have this mark. Communicative norms of the type articulated by Grice (1989), such as the cooperative norm, are fairly obviously of the formal variety. The mere fact that we would be communicating badly by violating the cooperative norm is not yet sufficient to generate a reason to be cooperative. We need to supplement these obligations with a claim about the value, desirability, or general reasons we have to obey these norms in order for these norms to imply reasons. The reasons to communicate well are, of course, nearly always present so we don't often need to explicitly articulate them. Moreover, given this fact about the nigh-omnipresence of general reasons to communicate well, it really is the case that saying "yeah, what I said had nothing to do with the topic of our conversation, but I don't care about that" seems really weird; to my ear as weird as the moral case. So the fact that it is typically unacceptable to respond to someone pointing out a moral obligation with "but I don't care about that" or "meh, morality" doesn't cleave between substantive and formal obligation since it is similarly unacceptable to respond to a conversational obligation with "but I don't care about that" or "meh, communicating well."

I do think, however, there is something substantially correct about the idea that the difference between authoritative and merely formal normativity has to do with when these obligations or, rather, the norms underlying them are IN FORCE. Unfortunately, this cannot be the entire answer to the story since rules of etiquette are also nearly always in force. Yet, at least in myth and legend, we still need additional reason to comply. Legal norms can also be in force without providing reasons. From the other direction, it is not at all obvious that moral norms are always in force; couldn't they be (in fact, aren't they...) suspended in various dramatic circumstances, so long as the exempting conditions were salient enough? And, yet, if that's enough, note that conversational norms are always in force *unless* we are in an explicitly articulated case in which they don't apply, such as speaking at a play or just talking nonsense. These are supposed to be merely formal though.

Even if we could resolve these problems, this type of answer is still unsatisfying. We would

also need an explanation of why moral rules are always in force, which itself demands some further feature of morality that underwrites this constancy. We will return to this sort of answer below. First, a few more substantial attempts at a characterization of the intuitively substantive.

2.1.3 Contentfulness

For our next attempt, note that there is no serious dispute that there *are* formal moralities, in the sense of moral rules propagated by this or that community. Some content of these seems widespread and authoritative, such as the typical fragment of local rules concerning not stabbing the occasional random passerby. Some content does not, such as religious- or honor-based permissions to hurt others. So the claim cannot be that the *subject matter* of morality is such that any widely-accepted system of obligations and permissions concerning it is robustly normative. Some systems of norms concerning this subject matter, such as honor-based local moral norms, seem to not be robustly normative. That is, unless we claim that the subject matter of local formal moralities is different than that of morality, but this strains belief. They concern the same thing, even if the local morality isn't necessarily correct.

Given that subject matter won't decide which systems of obligation are authoritative and which are not, we could try to use the fact that we can wonder about whether a particular local morality tracks big-M Morality, whereas we cannot really wonder whether a particular local etiquette tracks big-E Etiquette, to cleave these cases apart. Here, intuitively, what you see is more or less what you get. Likewise, there is no real question about whether local pool rules track big-PR Pool Rules; even more than big-E Etiquette, there is no such thing as big-PR Pool Rules. This is not to deny that we can wonder what the rules of small-e etiquette are; Emily Post exists for a reason, after all. The suggestion is that we can wonder what the *right* morality is, whereas we cannot wonder, sensibly, what the right etiquette is. Of course, we can wonder what the etiquette that would be most democratic, nicest, most fair, and so on are, but these are different questions. This is like wondering whether chess with or without *en passant* is a more fun or more elegant game.

The first stab at characterizing this fact builds on the idea that some merely formal obligations are built by fiat, whereas substantive morality seems not to be. The rules of pool are established by a group decision. There is thus no questioning them as the rules of pool, since the fact that they are the rules of pool seems to be justified completely by the fact that that's how we decided pool was to be. Likewise with later clarifications of pool rules for particular games. The law, similarly, seems to be justified by the fact that we have laid down the laws in a particular way, though the history and, perhaps, the conditions for what could count as a law at all are more complicated.

Unfortunately, this will not work since, again, there are cases of merely formal obligations which seem not to come into existence this way. Etiquette, for one very salient example, did not get stipulated into use by group decision the way pool rules did (at least in myth and legend). Likewise, prescriptive grammar norms (in the sense of how we *ought* generally to speak, not just fiddling pedantry about 'that' and 'which') did not get stipulated into being by fiat. So we need to approach our intuitive criterion a different way.

The more reasonable way of approaching the intuitive criterion takes seriously the idea that local moralities are meant to track big-M morality and, thus, that our local moralities can, as a body of norms, be either *correct* or *incorrect*. This tracks the fact that moral norms are intended

to have more descriptive content than norms of etiquette. With norms of etiquette, once we have found out what the local rules are, we have found out all there is to know, and there is no further questions about whether or not these rules are correct or not. Likewise, with pool rules, once we have read the board with the rules, we have found out more or less all there is to know (we also need the implicit unwritten norms, of course, but this doesn't affect the point). With morality, however, it seems like there is still more we can and typically do ask about. We want to know whether or not the moral rules correspond to the moral facts, whatever these are.

Of course, if this is to be more than a tablepound, we need to not simply presume that there are moral facts and that moral norms are meant to track them whereas there are no etiquette facts and etiquette norms aren't meant to track them. Our data is, rather, that it's weird and seems out of place to ask whether club rule etiquette is correct, whereas this doesn't seem out of place to ask whether Judeo-Christian morality is correct. Unfortunately, as soon as we make this clear, then there are other systems of norms that seem to have this character. Consider again prescriptive grammar. There are clearly rules for this—it is incorrect to use a singular verb form with a plural subject, as in “The aliens is coming”—but many of these rules are non-obvious and we quite often dispute them. Of course, there is also the interesting question of whether our prescriptive grammar—how we ought_{grammar} to talk—corresponds with our descriptive grammar—how we do talk. There is the question of what our prescriptive grammar should be like; inspection of disputes about typical pedantic bits of prescriptive grammar focus on whether or not the rules eliminate ambiguity, simplify our expressions, etc. Neither of these questions is about whether our prescriptive grammar is correct, but about what our prescriptive grammar ought to be like. Still, even in spite of this, there is a recognizable notion of a descriptively better or worse set of prescriptive grammatical rules and thus, plausibly, a notion of whether a prescriptive grammar can be incorrect in seemingly the same sense that a local morality can be incorrect—they may not correspond closely enough to Morality or Grammar.

Of course, this doesn't show that there aren't systems of prescriptive grammar that are equally “correct” in whatever sense perverse pedantry about splitting infinitives is incorrect. What it does show is that the simple fact that it makes little sense to wonder what the real etiquette norms are, whereas it makes perfect sense to wonder what the real moral norms are, doesn't do the work we want in pinning down the substantive systems of norms. Grammar, which gives rise to merely formal obligations when considered by itself, has the same property as morality. Interestingly, in the case of grammar, it is plausible that there equally correct systems of prescriptive grammatical rules. This leads into our next suggestion for a criterion cleaving the two apart: if morality is supposed to be *universal*, whereas a useful system of prescriptive grammar can take many shapes, then perhaps a combination of CONTENTFULNESS and some notion of universality will form a criterion for substantive obligations.

2.1.4 Universality

Morality has often been supposed to be universal in the sense that there is exactly one set of moral norms that is correct and, at least when applicable, in force for all agents. This is not to deny that local circumstances cannot modify the outputted directives of morality. Obviously if part of what it is to be moral is to not hurt each other and, given certain non-actual facts about air density and local plant life, throwing a fruit at someone doesn't actually cause them harm and is typically taken as an expression of good intentions, there will be no derivative rule forbidding me chucking my apple into an audience. In contrast, given the actual facts about apples and air density, I ought_M not throw my apple into the audience. And not only should I

not throw my apple into the audience, no one in my circumstances, regardless of which norms they personally subscribe to, should throw their apple into the audience.

The idea that there is a uniquely correct system of *pure* moral norms that applies to all agents is widespread and it does seem that the conjunction of CONTENTFULNESS and UNIVERSALITY pins down a notion of substantive obligation. It is, it turns out, quite difficult to find many other systems of norms that clearly possess both features. Perhaps the most obvious of the ones that do so is *epistemic norms*, especially when pitched in terms of norms like: avoid contradiction!, increase coherence!, and most obviously, respect the evidence! Of course, there are different appropriate methods of belief formation in the sciences, in ordinary day-to-day life, and in philosophy, but presumably this can be chalked up to the differences in the quality and types of evidence available these differing cases, just as differences in local morality can be chalked up to differences in our environment and sensibilities. Epistemic norms, the correct ones anyways, are supposed to be more or less universal. Epistemic norms also purport to meet our CONTENTFULNESS mark since, as decades of epistemology testify to, there is active and important debate about what the actually correct epistemic norms are.

This is all to the good since the nearest contender to morality as generator of substantive obligations is epistemic obligations. We thus have some confirmation that our joint criterion is on the right track. We get further confirmation by noting that one of our previous criteria, IN FORCE which had some bite is naturally subsumed under universality. Insofar as the correct moral norms are supposed to hold for all agents, they will plausibly be IN FORCE, whenever they are IN FORCE, no matter which agents we are considering. We still have the problem of explaining why moral norms are in force when they are, but presumably an explanation of why they are universal will go on to furnish an explanation of why they are in force, when they are in force, as well.

2.1.5 Recap and Upshot

It looks as if the best mark of authority, at least from among the ones I have surveyed, will be something like a combination of CONTENTFULNESS and UNIVERSALITY. GRAMMATICALITY and DIRECTIVENESS are non-starters, BUT I DON'T CARE ABOUT THAT seems to cut too broadly, and IN FORCE can be subsumed under UNIVERSALITY. For the rest of this essay, we will assume that this combination partially characterizes the mark of authority for us. So insofar as we take a system of norms to be substantive, we take it to both be an attempt to describe a set of real facts about how we should interact with the target of that set of norms—for moral norms, actions, for epistemic norms, belief formation—and that substantive sets of norms will be universal in the sense of holding for all agents.¹³ There may be other ways of characterizing types of normative authority, but this seems to capture most of the sense in which moral and epistemic norms are said to be substantive whereas etiquette and rules of pool are not.¹⁴

I say “almost” because it does not yet follow that a system of norms that satisfies these marks induces a notion of obligation which is anything like reason-implying in the sense de-

¹³Note that narrowly described systems of norms which are *constitutive* of a game—such as the rules of tic-tac-toe—may be trivially universal since anyone not playing by these rules is engaged in a different activity. Such norms will not be contentful. Thanks to Derek Baker for useful discussion.

¹⁴See, for example, McPherson (this volume) for an attempt that uses settling disagreements between conflicts of norms as the identifying mark of at least one substantive bit of normativity, the practical ought of deliberation. It isn't clear to me whether McPherson takes this ought to be the only substantive bit of normativity. Since it isn't the only common-concept set of substantive norms (morality and epistemic norms being also common-concept pretenders to substantivity), this sort of analysis won't work for my purposes.

scribed above. For all we have said, it could be that a system of norms is correct for all agents and in force, inducing an “substantive ought”, without this implying that we have any reason to comply with it; this would mean that it has all the marks of substantivity without being substantive in the sense defined above of implying reasons simpliciter to comply. To see this point, consider legal norms. These norms are not universal, but put that to the side for now; presumably, when one is present in a country, these norms are in force for you (modulo certain complications, but put those aside). They give rise to a recognizable notion of legal obligation and to a subscribed legal ought. Yet, it seems, they do not automatically have reason-providing force.

As Hart noticed long ago, there is a gap here to be filled since we can understand what it is to be in conformity with your legal obligations without yet providing an answer to why we have any reason to comply with our legal obligations (Hart 1961). If, say, we could care less about punishments so-incurred, and where violation wouldn’t disenfranchise others, and so on, it is difficult to see what reason we have to comply. Legal norms do not meet our mark of the authoritative, so we do not have a case of norms which meet the mark and which are not reason-implicating.¹⁵ Given the worries just laid about for legal norms though, it does seem a substantive philosophical hypothesis that the marks of authority so laid out suffice to generate reasons. We do, however, take as one of the characteristic features of a substantive obligation—at least as I have so far characterized it—that it generates reasons. So we need to add to our mark of authority something which bridges this gap.

I think the right suggestion for what to add was given at the beginning. Remember that one of the characterizations of normative authority had to do with the idea of guidance. Substantive norms are meant to guide our practical deliberation, irrespective of our general interests, unlike etiquette norms or game norms. Legal norms, like etiquette norms and game norms, only guide when we have an interest in obeying the law, unlike moral or epistemic norms. So guidance looks like just the thing to carve apart legal norms from substantive norms and close the gap in our intuitive notion of the substantive.

Guidance, along with UNIVERSALITY, will suffice for some rough idea that these norms are reason-implicating and separate out the sense in which legal norms are in force from the way moral or epistemic norms are in force. If, that is, norms X are taken to be both guiding and universal, then insofar as they are in force, we should take them to bear on our deliberations about what to do,¹⁶ and this thereby plausibly give rise to the idea that we have *pro tanto* reason simpliciter to do as we ought_X.¹⁷ Of course, what we take to be reason-providing and what really is reason-providing might come apart, but at least this gives us a plausible mark of the normative.

The reader might wonder at this point why we don’t just take reason-implicating as our mark of substantiveness in the first place. The reason is that reasons, in the sense used in the real definition of substantive given above, are a philosopher’s construct. Reasons, in the philosopher’s sense used above, are really normative, not merely the sort of thing that we are inclined to cite in favor of acting such and so ways. And this is a good thing since in ordinary parlance, we give “reasons” to comply with social norms just by pointing at the norms themselves.¹⁸ It is,

¹⁵Though see Kolodny (2005) and Maguire and Woods (ms) for the claim that epistemic norms and rational norms are also not *eo ipso* reason-implicating.

¹⁶That is, modulo the point above that explicitly considering them might be problematic for Railtonian reasons.

¹⁷Note that I am not suggesting that all reasons will guide us in this sense. My claim is that norms which should guide us, which are in force, and which are universal are plausibly reason-implicating.

¹⁸See Heath (1997), Tiffany (2007, §2), and Joyce (2001) for useful discussion of this point.

however, a significant philosophical step to the claim that these norms imply real reasons and, for the obviously merely formal norms, a fairly dubious one. Guidance is something that is not a philosopher's construct, as anyone who was ever taught social norms at their parents' knees knows. We were nearly all taught that when figuring out what to do, we should think about what's right and wrong, after all.

From guidance we can infer, by means of UNIVERSALITY and CONTENTFULNESS, that an authoritative obligation is taken to be really reason-implicating in the philosopher's sense without attributing to ordinary folks the philosophers' notion of a reason. "Taken to be", not "is". It is a substantive philosophical step (here denied) from the claim that we take an obligation to be authoritative to the claim that it is, in fact, authoritative in the sense from.

2.2 Merely Formal Obligations

Having delineated a proposed intuitive mark of substantive obligation, we now turn to examining merely formal obligations. This is important since there is a tendency to shrug off merely formal obligations. Merely formal obligations really are obligations though. Shrugging them off is a serious mistake, and all the examples of ridiculous norms that would never be actually in force should not convince us otherwise. Seeing how merely formal obligations are, in a very important sense, obligations helps to motivate my account below as a convincing explanation of the real difference between substantive and merely formal obligation. We want a characterization of the substantive that (a) keeps in place that substantive obligations are different than merely formal obligations, but which (b) keep in mind that systems of norms which give rise to substantive and merely formal obligations are substantially similar in that they both give rise to real honest-to-god obligations.

2.2.1 The Obligatoriness of Obligations

The first point to emphasize is that some merely formal obligations, in fact, all of them under special circumstances, really are obligations. This point is often elided over by working with examples of systems of norms which are not and, sometimes, even could not be in force. Norms which establish a set of correctness conditions and, thus, a subscripted obligation are cheap after all. When we turn to merely formal systems of norms which are actually in play, we find a striking similarity between the obligations we find there and the obligations we find in substantive systems of norms.¹⁹ So let us look a bit closer at systems of norms that really are often in force: promissory, game, and etiquette norms.

Consider first promissory norms. If we abstract away from overly moralistic views of promissory obligations, the resulting system of norms allows us to make promises to do things which we pretty clearly ought not, all things considered, do. We might promise to do away with the person who cut our friend in line, we might promise to pursue our mother's campaign of total world domination. Are these promises obligatory? It seems so. We seem obliged, in at least the subscripted sense, to keep immoral promises. If we fail to do so, some criticism on the behalf of the promisee is legitimated (even if they ought, all things considered, not actually criticize us.)²⁰

¹⁹Mitchell Berman (forthcoming) calls this similarity in normative structure between different systems of norms a *normative isomorphism*. Berman correctly points out that this similarity holds not only between cases like legality and morality, but also between game normativity and morality. The point could also be made, as I made it above, in terms of the linguistic semantics for the 'ought' which is given rise to by means of these systems of norms. Full development of this linguistic point would take more space than I have given it here.

²⁰See (Woods 2016) for significant further discussion of these sorts of cases.

For example, consider the case of Omar Little's testimony against Bird in the Wire. Omar has not seen Bird shoot Gant, which the police know full well. Nevertheless, they let him take the stand and perjure himself. Bird is a dangerous criminal and definitely worth being put away. Moreover, given Omar's dispositions to rob only criminals, putting him away would be all things considered a bad thing. Nevertheless, I submit that Omar ought_{legal} to have not perjured himself and Stringer Bell's complaint that he had to be warranted. Of course, given that morality, prudence, and justice all favor Omar perjuring himself, we ought not to criticize him for doing so. In my terminology, he is not blameworthy for perjury, though he is certainly blame-liable for perjury. For another example, suppose that I have promised to do something immoral in a particular way. I do the selfsame immoral thing, but in a way different than I have promised. I submit that I am blameworthy for both doing the immoral thing and, moreover, for doing the immoral thing in the wrong way. The person to whom I have promised to do the immoral thing is legitimated in criticizing me for doing both things, though, of course, that person is herself criticizable for accepting my promise in the first place.

Consider next obligations of etiquette. Our etiquette obligations pretty clearly don't lapse even when they're ridiculous. Even more so, they don't lapse since we are indifferent to them. As Foot put the point, both morality and etiquette

...are inescapable in that behavior does not cease to offend against either morality or etiquette because the agent is indifferent to their purposes and to the disapproval he will incur by outing them. (Foot 1972, p. 311)

For a more concrete example, take the norm that we should not wear a hat at the dinner table or that we should wear shoes when we are in a restaurant. There are contexts in which doing the former or failing to do the latter are signs of disrespect, but even when no one would take it as such, it is still a violation of etiquette to do so. People are then within their rights to criticize someone for violating these norms.

Of course, both promissory and etiquette norms are flexible in the sense that they can be lifted by appropriate authority figures in local contexts. A restaurant can allow, explicitly, bare feet and the head of household (or heads of household) can lift the ban on wearing a hat at the dinner table. Nevertheless, when these norms are not relaxed—when they are *in force*—then they really do furnish obligations, at least of the subscripted variety. And, in most natural cases of subscripted obligations, they permit a certain form for thin blame to be leveled at the violator. We can correct and tut-tut the person sitting at the table with a hat on, we can cluck our tongues at the person who strolls into a restaurant without shoes on. We can complain about the person who jumps a ball incorrectly in pool or refuses to let us take a pawn *en passant* while knowing full well that it is permitted in the game of chess being played.

This is true even when there are reasons to do things differently. Even if the hat rule is a remnant of some bizarre religious custom, even if the shoe demand infringes on the rights of those who don't want to wear foot-prisons, even if pool should allow bottom-jumps, even if chess with *en passant* is an inferior game, given that these are the norms as they are, we are obliged to respect them.²¹ Silly games, etiquette which is relic of an earlier time, and dressing standards still have rules and, insofar as we are playing them, we are obliged to play as they are supposed to be played and we are culpable for not so playing when we don't. This, however, might not be enough to give us reason to play as we ought.

²¹See (Maguire and Woods, ms) for discussion of this point and responses to objections to it.

Consider norms governing correct play in a game of chess. We are obliged, when playing chess, to not move our rook along a diagonal. We are permitted to move our rook in a straight line as far as we like. Suppose we were to be playing chess with our nephew who was, alas, susceptible to crying jags when they lost and who didn't take losing as a lesson in how to play better. Moreover, suppose we could care less about the game; we are just passing time until leaving for a movie. It would be all to the good, given this situation, if we didn't criticize our nephew for playing moving their rook along a diagonal. Likewise, it would be all to the good, given this situation, if we convinced our nephew that rooks couldn't move all the way down the board in a straight line (if doing so would bring about an immediate loss for them). Our nephew violates his chess obligations when he plays that way and we are in a position to criticize them for doing so, even though so-criticizing is probably a bad idea. We violate our chess obligations when we convince our nephew to play incorrectly and they (or a bystander) is warranted in criticizing us for doing so, even though it is probably a bad idea to do so.

Contrast the cases just described with cases of systems of norms not currently in force. Suppose we see people moving chess pieces around on a board and we step in to politely correct some seemingly illicit play. The players are entitled to respond that they're not playing chess and this would be a complete defense against our criticism. Likewise, suppose our tedious uncle started upbraiding us for wearing a hat to the dinner table. If our mother steps in and says that it's allowed, then that's a complete defense against his complaints, given the norms of etiquette as they are.²² Likewise, it's a complete defense against someone not getting in contact after saying they would when it's pointed out that "we really must hang out", in the United States, does not imply any commitment to hanging out (in fact, quite the opposite.)

In any of the cases where a system of norms is in force, it would be completely bizarre to respond to criticism with the claim that you are not subject to them. We can say "I don't care about being polite", of course, but there is still a residue of strangeness to doing so when politeness norms are clearly in effect; it really is warranted for others to criticize us on the basis of our violation of politeness norms. Whether or not that gives us reason to be polite, of course, is still another matter. If we will shrug off the criticism and there is no reason to bow to social pressure (or, in fact, reason to buck social pressure), then we may not have any reason be polite. Still, we cannot object when someone points out that we have acted incorrectly.

A sensible account of normative authority should preserve this sense of genuine obligation which occurs when a system of norms is in force. It would be bizarre to think that we are entitled to buck our in force obligations, be they chess obligations, etiquette obligations, or even promissory obligations on the grounds that they are merely formal obligation. Of course, we now need to make good on what it is for a system of norms to be in force. We turn now to this.

2.2.2 When are Norms In Force?

There seem to be two rough types of systems of norms that give rise to merely formal obligations, dividing exactly over when the norms are in force or not. The first type are systems of norms that hold of a situation by means of some explicit or tacit agreement of a group of people. Game norms seem to be the paradigmatic example of such norms; we are obliged to follow the rules of chess—or at least the variant agreed upon—because we have agreed to play chess. We are

²²This will not work for some norms for reasons set out below. We do not always have the option of relaxing the norms in play.

obliged to play local bar rules pool because we have voluntarily started playing pool in a bar in such and so location. We can opt out of being subjected to these norms by simply not playing, though while we are playing we are subject to them.

This, however, isn't quite right since which rules we play a game by aren't often entirely voluntarily. Given that we are playing *chess*, not *schmess*, we are bound to play by means of some recognizable version of chess rules. Likewise with pool and any other game. So the real divide is between what we might call boutique merely formal norms that are explicitly or tacitly agreed upon and those where we can engage in the activity governed by merely formal norms, but the shape of these norms isn't itself entirely a matter of agreement. The former type of merely formal norm isn't especially relevant for our purposes, so we will put it aside as a special case of the merely formal norms where the activity is more or less voluntarily engaged in.

The other type of norms which give rise to merely formal obligations concern activities our engagement with which is less voluntarily and where the shape of the norms is less obviously a matter of tacit or explicit agreement. Etiquette, for example, is inculcated into us from birth by our family, friends, and social environment. Likewise, norms of prescriptive grammar are not voluntarily agreed upon; we pick these up as we learn to speak. This is not to say these norms are *always* in force. There are certainly conditions under which norms governing grammatical correctness lapse, such as writing poetry. It is likewise very plausible that under certain conditions, such as being stranded on a desert island, norms of etiquette lapse. However, these conditions are excusing conditions, the default is to be bound by the merely formal obligations arising from grammatical and etiquette norms. These norms also seem to apply more broadly, being largely in force, as opposed to the first rough type of norms we characterized above.

What explains the difference between these types of merely formal norms? The natural thing to point to is a combination of two facts. First, that these norms have a certain form of involuntary currency in the context at hand. American etiquette norms are in force in America because these are the tacit norms we generally internalize as we grow up or dwell in America. This point is banal, but it is important. Norms are partially in force because they have currency for us and they have currency for us either because we have agreed that they will or because they have been internalized (typically involuntarily) by us as the norms governing a particular activity. Second, these norms are targeted at certain characteristic activities that are a basic part of human interaction and that we typically do not engage in these activities due merely to a voluntarily choice. We typically do not get roped—in the requisite sense—into playing a game of pool. We are quite often roped into communicating in our native language or interacting with others.

Interestingly, it is often because these norms interact that we are involuntarily put under their obligations. For example, suppose someone pops their head in my office and asks me if I know where Freddy's office is. I could say nothing, but that would violate my etiquette obligations. I could instead tell them something, satisfying my obligations of etiquette, but then I would put myself under obligations to speak correctly. Situations like this are pervasive. Still, the point is that we often fall under merely formal obligations by means of the fact that we are engaged in a particular activity that they govern. Given this, we can mark a difference between types of merely formal norms by noticing how voluntary our engagement with these activities is and whether these norms have currency for us because of contractual agreement or because we have internalized them.

The merely formal norms of the first type are often cited cases for norms which are not substantive and there is good reason for this; we have no reason to comply with these norms unless we are engaged in the activity they govern and we need not do this. This should not obscure, though, the sense in which we are, in fact, bound by merely formal norms when we are engaged in their target activity. We do not get out of our obligations to not move our bishop in a straight line merely because we didn't want to play chess in the first place, or couldn't care less about chess norms, or what have you. We're playing chess and, when playing chess, chess norms are in force. When we break these norms, we are criticizable on this basis.

2.2.3 The Internal Point of View and Fetishism

Of course, we need to understand what explains why in force norms give rise to real obligations and how internalization works, but I believe the answer to that question was given to us long ago by Hart (1961).²³ Hart writes, concerning the "internal point of view":

[F]or it is possible to be concerned with the rules, either merely as an observer who does not himself accept them, or as a member of the group which accepts and uses them as guides to conduct. We may call these respectively the 'external' and the 'internal points of view.' (Hart 1961, pg. 89)

This is strongly suggestive of the idea that the internal point of view involves taking a system of rules to be reason-implicating or, anyways, to guide our practical deliberations. But this is not all Hart says. He also suggests, earlier, that:

For the judge, in punishing, takes the rule as his guide and the breach of the rule as his *reason* and *justification* for punishing the offender. (Hart 1961, pg. 11)

and this is suggestive of the idea of the internal point of view involves taking a system of rules to really be *obligations* in the sense suggested above (§0.3.1).

In earlier work, I suggested that Hart is glossing over (or, anyways, neglecting) an important distinction.²⁴ For us to take the internal point of view to a system of legal norms is for us to treat violation of the standard they induce—failure to meet our subscribed legal obligations—as sufficient warrant not to punish the offender, but to treat the offender as being *liable* for punishment. Going on use their liability for punishment as justification for the worthiness of punishment is a further step. To put it bluntly, the mere fact that someone is liable for punishment isn't always sufficient for them to be punishment-worthy. This distinction generalizes to all normative systems that build in a notion of criticism which attaches to violation of our obligations, as is the case for all systems of norms adumbrated so far, and is the case for most systems of norms period. This will become important below, but for now it suffices to mention the distinction.

²³See (Shapiro 2011, pp. 95-98), for a useful summary discussion.

²⁴As Derek Baker reminded me, Hart also makes a related distinction between the justification of the practice of punishment—in terms of a social good being realized by the practice—and the justification from within the practice of punishing someone—in terms of retributive considerations (Hart 1960). However, even if this takes us some ways towards my distinction, it isn't enough. Hart think that we can override the internal justification in terms of the social good being realized significantly more strongly. This, for instance, legitimates hanging someone out to dry even though it is known to the authorities that they did not commit the crime. I think, in these cases, it is more accurate to say that there is reason simpliciter to hang them out to dry, but that they are not-punishment-liable or punishment-worthy. I also think, contra Hart, that we do not need to justify legal systems in order for them to provide reason simpliciter. They provide reasons because we care about being legally-fitting or avoiding legal-sanction-liability.

So, in Hartian terms, we can gloss internalization of a particular set of norms d in terms of the tendency for a social group to take the internal point of view towards d . In particular, norms are in force when they are internalized by a social group and this internalization is the general tendency to regard the subscribed obligations grounded by this system of norms to be real obligations and to treat ourselves as liable to sanction because of our behavior with regard to these obligations. I call a combination of this view with a desire-based source of reasons simpliciter and the a picture of the properties grounded by it (see §3.2 below) QUASI-CONVENTIONALISM about normativity (see (Woods 2016) for significantly more discussion.)

So far, we have noted that some norms give rise to real, even if merely formal, obligations and that these norms are in force when we are engaged in the activity that they concern. Our final point to be made about merely formal norms is the possibility of what is sometimes called fetishism about norms. We are fetishistic about norms when we view norms, even merely formal norms, to be such that conformity to the norms—being *d-fitting* for some system of norms d —seems like something we have intrinsic (or, anyways, non-instrumental) reason to do. This phenomena should be familiar to anyone from reflection on their early experiences questioning increasingly truculent adults. It is the natural interpretation of what these adults mean by their incredibly frustrating response “look, that’s the way (to/we) do it!” to questions about why we should to let our aunt kiss us or why we need to clean up our room before the cleaner comes. It is the impression that the mere fact that this is how we ought_X do something means that we have reason—indeed, often compelling reason—to do it that way.

Some of us find that some intuitively merely formal norms sometimes self-present as having this feature.²⁵ I do not want to speculate (yet, but see below) about how our concepts came to sometimes have this feature.²⁶ The important thing is that this phenomenon really is phenomenon and it is one that many of us recognize in how we treat these norms.²⁷ Normative fetishism is often taken to be a bad thing, but I think this is an overreaction. We go wrong in treating silly systems of norms as fetishistic, of course, and sometimes our treatment of otherwise reasonable norms can be fetishistic in a criticizable sense, but it isn’t by itself a bad thing to *treat* a system of norms as the sort of thing that gives rise to reasons by itself, even a merely formal system of norms like etiquette or grammar. In fact, I reckon that it better not be a bad thing given how often we seem to do it. The important thing to note about fetishism is that when a system of norms is approached fetishistically by us, then it will feel somewhat substantive from the inside. It will *not* necessarily feel universal or contentful, but it will satisfy at least the guidance mark. This will become important below.

Recapping briefly, I have argued that merely formal norms are often in force and not by voluntary agreement between participants to the system of norms. Rather, these norms are often in force in light of the fact that we are engaged in the activity they govern, such as playing pool, speaking, or interacting with other people in our community. When a system of norms is in force, then it gives rise to real honest-to-god obligations which typically is attached to some form of blame or punishment for having behavior out of step with the norms. This might range from something like social sanction (for violating norms of etiquette or some games), actual

²⁵I actually believe this phenomena is ubiquitous, but given how reluctant people are to admit this kind of thing, I offer the hedged some/some/some formulation here so as to not offend.

²⁶One suspects that it is partially the role of early pressure from truculent adults for us to accept these norms as binding norms for how to behave and the tendency to close interpersonal deliberative discussion by citing these norms. Of course.

²⁷Exegesis isn’t the point of this piece, but I think Hart (1961) wavers between an interpretation of the internal point of view as being fetishistic and it merely supporting a notion of real obligation which is not necessarily taken to imply reasons in the sense laid out above.

punishment (for criminal behavior), or some internal punishment like losing the game being played. Importantly, when a system of norms is in force and attached to a form of blame or punishment, then we get a property of blame- or punishment-liability, different from blame- or punishment-worthiness, which attaches to violating the subscribed obligations arising from the system of norms. Finally, I noted that it is possible to be fetishistic with respect to a system of norms and that when we are d -fetishistic, for d a system of norms, then from the inside these norms and the obligations they induce will feel close to substantive.

3 Reconceiving the Distinction

Our main task in this essay is to reconceptualize the distinction between substantive and merely formal norms. To that end, I started by laying out what I wanted to deny: that that we had reason simpliciter to do something just because we ought $_X$ to do it, for ‘ought $_X$ ’ a subscribed ought. I then explored what we intuitively had in mind when we called a system of norms and the obligations grounded thereby substantive, settling on the idea that our intuitive notion of a substantive obligation involved the idea that the norms governing it were UNIVERSAL, CONTENTFUL, and that they provided GUIDANCE for our practical deliberations.

We then turned to merely formal norms, where I argued that merely formal norms often really did give rise to honest-to-god obligations. They do so when the norms are IN FORCE. Norms are in force when (a) our community has internalized them and (b) we are engaged in the activity which they govern. Norms which are in force give rise to subscribed, though genuine, obligations that we can be in compliance with or in violation of. When we are in violation of them, we will incur liability to sanction (if there is a sanction attached to violating our obligations in the relevant system of norms). Finally, I noted that we *can* treat systems of norms and the obligations they induce fetishistically, taking ourselves or others to have reason to do what we ought $_d$ just because we ought $_d$ to do these things.

With these pieces in place, we can articulate desiderata for our reconceived distinction between substantive and merely formal normativity. We want an account of this distinction which denies that there really are substantive obligations in the philosophical sense defined above, but where the actual distinction:

- allows that there really is an important distinction between substantive and merely formal obligations. (§2.1)
- respects the sense in which merely formal systems of norms and the obligations they induce are really obligations (§2.2.1)...

...without requiring us to treat the ‘ought’ of merely formal obligations and the ‘ought’ of seemingly substantive obligations as ambiguous (§1.1)

- captures, as much as possible, the intuitive characterization of substantive obligations we dredged up above (§2)...

...though it is no requirement that the theory vindicate the content of this characterization in the sense of substantive obligations really being, say, contentful. Rather, we need to explain why this seems to be a property of substantive norms and, ideally, why CONTENTFULNESS tracks something of importance.

- allows that disputes between obligations, even those of different types, can often be arbitrated (§1.1)...

...without implying that all such disputes can be arbitrated. (§1.1)

We can, it turns out, meet all of these desiderata. In order to do so, we need our background source of reasons to obey two principles:

DESIRE-BASED-REASONS That I *want* to do something or take something as an end is or gives rise to a *pro tanto* reason (simpliciter) to do it.

MINIMAL INSTRUMENTALISM If I have reasons to do something and something else is a necessary means to do that, then I have reason to do that something else.

But both of these are supported by any number of background pictures about reason simpliciter.²⁸ I will briefly sketch the solution here, then turn to defending pieces of it below.

3.1 The Account

Given a system of norms d containing forbiddings, permissions, favorings and disfavorings, we presume that d induces a subscripted notion of ‘ought $_d$ ’. Let a person be d -fitting when are in conformity with their d -obligations—when they do as they ought $_d$. To act in a morally fitting manner is do as we ought $_{morally}$, to act in a etiquette-fitting manner is to do as we ought $_{etiquette}$, to act in a legally fitting manner is to do as we ought $_{legally}$, to act in a promising-fitting manner is to do as we ought $_{promising}$. Of course, we do not generally describe these properties this way. Rather, we have lexicalized some of them under different terms: good, polite, law-abiding, and possessing integrity. For more game-oriented systems of norms, we talk about being sportsman-like, fair, and so on. These terms play a role in making accessible the rather formal property of being in conformity with the obligations of a system of standards.

Suppose d is a system of norms which (a) has currency for us and (b) is sometimes in force. Then we will typically have some lexicalized adjective to describe what it is to be d -fitting. Importantly, the property picked out by this adjective can be the content of desires and ends. We can desire to be (morally) *good*, *law-abiding*, *polite*, and *to have integrity* or to be *reliable*. These desires themselves can be instrumental or intrinsic. We may desire to be law-abiding because so doing will keep us out of trouble. We may desire to be law-abiding merely because.

When d also contains a system of sanctions for failing to be d -fitting, then we have another salient property, being *sanction-liable* (Woods 2016). Sanction-liability can also be the content of our desires and, again, can be such for both instrumental and intrinsic reasons. d -sanction-liability and d -fittingness are related, so some of our desires will also be so related. We might desire instrumentally to have d -fitting behavior *because* it will let us avoid d -sanction liability (in fact, it’s the only way to do so if d is in force). Given MINIMAL INSTRUMENTALISM, we will then have a derivative reason to have d -fitting behavior. Most of us, most of the time, will desire to avoid things like liability to moral sanction (moral blame). Many of us, most of the time, will desire to be morally good or to have morally-fitting behavior. Given DESIRE-BASED-REASONS, this means that most of us, most of the time, will have reason simpliciter to act as we ought $_{morally}$.

What characterizes substantive systems of norms? First, we will restrict our attention to norms which are nearly always in force. Substantive systems of norms tend to govern activities

²⁸It is likely we can actually replace DESIRE-BASED-REASONS with something similar which does not route through desires, but a number of the contours of the view I will describe below would have to be rejiggered in implausible ways. I will thus presume that any reasonable account of reason simpliciter attributes some weight to the things we desire, even when this weight is often outweighed by competing reasons. I have never seen any even moderately plausible reason to reject this claim.

which we are engaged in nearly all the time. Now, if we add a substantive empirical premise, that nearly all of us tend to strongly desire to be d -fitting when d is a stereotypical substantive system of norms like morality, epistemic norms, and the like, then we can distinguish between substantive and merely formal norms. Substantive systems of norms, in addition to governing activities we are nearly always engaged in, are those systems of norms d where there is a widespread tendency to intrinsically desire to have d -fitting behavior. An intermediate case, which I will call *substantivish* norms, are those systems of norms d , governing activities we are nearly always engaged in, where there is a widespread tendency to intrinsically desire to avoid d -sanction-liability (and thus, derivatively, to have instrumental desire to have d -fitting behavior). Finally, merely formal norms are those norms d , governing activities we are nearly always engaged in, where there is at best widespread instrumental desire to avoid d -sanction-liability.

I opened by claiming that no system of norms and the obligations grounded thereby were substantive in the sense of being sufficient unto themselves to imply reasons simpliciter to comply with them. However, given that we typically have intrinsic desires to either be d -fitting or, alternatively, instrumental desires to avoid d -sanction-liability by being d -fitting, we will typically have reasons simpliciter to do what we ought_{morally}. We can then define a local d -specific notion of a reason thus:

- A d -reason to ϕ is an instrumental reason simpliciter to ϕ whose weight derives from our (desire-based) reason to have d -fitting behavior.

For example, epistemic reason to believe that I'm writing this essay derives from my reason to have epistemically-fitting behavior (by DESIRE-BASED-REASONS); its weight comes from the strength of my desire to have epistemically-fitting doxastic states. I have strong evidence that I'm writing this abstract, so not believing that I'm writing this paper would be unfitting. So, by MINIMAL INSTRUMENTALISM, insofar as I desire to have epistemically-fitting doxastic states, I have reason to believe I'm writing this abstract.²⁹

We can now recast our distinction in more palatable terms. A substantive system of norms d is one where we typically have d -reason to do what we ought _{d} . A merely formal system of norms d is one where we do not typically have d -reason to do what we ought _{d} . Substantive systems of norms can be subdivided, as I did above, into substantivish and substantive systems of norms, but I will not do so for the rest of this essay.

The reader may have noticed that I presumed that the systems of norms we were distinguishing were those which were typically in force. The reason for this is that I think the in force/not in force distinction has muddied the waters in distinguishing substantive from merely formal norms. Even substantive norms can not be in force—consider, possibly, morality in catastrophe—and merely formal norms can typically be in force—consider basic norms of etiquette. So I will restrict myself to norms that are typically in force from here on.³⁰

²⁹Note that this defined notion of a d -reason differs from what we might see as a moral reason from “inside” morality, so to speak. From the point of view of morality, for example, we might have decisive reason to not play a nasty joke on Fred. But, given that we might desire the aesthetic joy of playing a successful nasty joke on Fred and care only a bit about morality, the weights of these two reasons will come apart. Moreover, internal moral reasons are not, by themselves, real reasons; they only approach real reasons when they are backed by the appropriate desire to be morally fitting.

³⁰The reader might worry that this means that things like chess normativity are substantive. This, I reckon, is nearly correct. They are substantivish since we typically, though not universally, desire only to not be liable for having game-unfitting behavior. They can be fetishized, of course, and sometimes are, and when they are they should be treated as substantive norms *when they are in force*. This point is not without precedent, Berman (forthcoming) sensibly lumps them in with substantive reasons and I follow him.

We're now nearly in a position to capture much of the intuitive characterization of substantive obligation. If we strongly desire to be morally fitting, then we should let moral norms play a role in our practical deliberations since we will often have actual reason simpliciter to comply with our moral obligations.³¹ Substantive norms are typically always in force, so we will typically have strong reason to comply with them. What we are missing is the contentfulness and universality of substantive norms. Substantive norms, characterized as I have above, are clearly still hostage to both context and our desires. Below, I will argue that even though this is true, we can explain why we see substantive norms as contentful and universal.

Our account matches up nicely with the rest of the desiderata. The obligations arising from formal systems of norms and substantive systems of norms are grounded in the same way, through the fact that the systems of norms have currency and are in force. The reasons which arise from our desires to act in accords with our subscribed obligations are likewise explained in the same way. The major difference will be that we do not typically desire intrinsically to be *d*-fitting or to avoid *d*-sanction-liability when *d* is a merely formal system of norms. We do, typically, desire both when *d* is a substantive system of norms. Given these similarities, it is no surprise that our reasons to have *d*- and *e*-fitting behavior are often commensurable; they are same type of reasons and naturally can be weighed against each other. It is similarly not surprising that we sometimes incommensurable reasons to act *d*- and *e*-fittingly; desires are very coarse grained and often when we put them on a scale they're roughly on a par. Given how nicely it captures these desiderata, this account goes a long way in capturing the distinction between substantive and merely formal obligations, so long as we can close the remaining gap between our picture and the intuitive characterization of substantive norms laid out above. We will turn to this point shortly, after a brief discussion of (a) the relationship of systems of norms to obligations, (b) the relationship between *d*-sanction-liability and *d*-fittingness, and (c) the role of the abstract notion of *d*-fittingness as opposed to the actual prescriptions of a system of norms.

3.2 Quasi-Conventionalism

QUASI-CONVENTIONALISM is the view that systems of norms, which have currency and are in force for us, give rise to subscribed obligations in the manner suggested by Hart (see §2.2.3 above), ground the property of fitting these obligations, and the property of being liable to sanction. I have called these latter two properties *d*-fittingness and *D*-sanction-liability above. These properties figure into our conative states in a broader and more general way than the actual notions of what we ought_{*d*} to do. They are a social phenomena as we need not subscribe to a system of laws for us to have subscribed obligations, for us to be evaluated in terms of our fittingness with these obligations, and with regard to the applicability of these sanctions to us.

What do I mean by this? Roughly, I just mean that for some systems of norms, we engage in them merely by our participation in social environment. They are, so to speak, the cost of doing business. Politeness norms in a foreign country, for example, apply to us even if we think they're silly (such as taking my shoes off in someone's house). Likewise, disagreeing with the practice of promising doesn't get us out of our promissory obligations, though it might mean that we have no reason to keep them. And, of course, we are entitled to upbraid someone who has been in the states for a few months, but hasn't started tipping—they are bound by our fiscal norms, even if they think they privilege individual excellence over egalitarian values. Mr. Pink's fevered explanation of his view of tipping, for example, causes consternation exactly because

³¹Though, again, not necessarily by reflecting on what moral norms ask us to do since this might lead us astray. How we should reflect on and deliberate with moral norms will depend on the contours of the moral norms.

he's bucking the obligations he has, and the critical response of the rest of the crew in Reservoir Dogs is exactly what he has coming.

There is a lot of detail to be spelled out about how this structure works, but again I'm going to foist most of this work off on Hart and other social theorists. It seems eminently clear to me that we need these notions since, whatever else one thinks, a similar story is obviously correct about our actual linguistic norms, scientific norms, game norms, and the like. These norms create formal obligations because they have currency for us; they have currency for us since we take them to have currency for us; we criticize others and apply sanction in light of this fact. We now turn to the relationship between *D*-fittingness, *D*-sanction-liability, and *D*-sanctionworthiness.

3.2.1 *D*-Fittingness v. *D*-Sanction-Liability v. *D*-Sanctionworthiness

It is important to distinguish the three subsubtitular properties here. *D*-fittingness is a formal property attached to being in accords with the findings of a system of norms. We are *d*-fitting just in case we do what we ought_{*d*}, we are *d*-unfitting otherwise. Typically, when we are *d*-unfitting, this will mean we are liable to sanction, but not always. This will depend on the conditions for sanction built into the system of norms and, of course, whether the system of norms has a built-in sanction at all. I have presumed that substantive systems of norms typically have a built-in norms, but this is merely a simplification.³²

Sanction-liability, though, is to be strongly distinguished from sanction-worthiness. It is easy to worry about whether the quasi-conventionalist picture of normativity is correct if these two are confused. It is rather *implausible* that someone is blameworthy, say, merely on the basis of violating social norms in an innocent manner. It is incredibly plausible that someone is blame-liable for doing so. Compare the distinction between prosecution-liability and prosecution-worthiness. If, say, I wander over a do-not-cross line, but do so in a way which isn't malicious, but stupid, then presumably I am liable for prosecution, but I'm not prosecution-worthy.³³

In my earlier work, I suggested that our promissory reasons were sourced in our desires to avoid blame-liability for promise-breaking. Though I still think this is typically the correct analysis for our promissory reasons, further reflection (Maguire and Woods, ms) has inclined me to take a more general stand here. I think, in fact, that our *d*-reasons can be sourced in *either* our desires to avoid sanction-liability or our desire to be *d*-fitting. In both cases, we end up with reason to be *d*-fitting, though this desire might be instrumental or intrinsic, and then derivatively with instrumental desire to do as we ought_{*d*}. We can go on to make use of this further difference to distinguish what I called above substantive and substantivish systems of norms, but I will not do so here. Either way, we have the sort of reasons we need to conform to what we have subscribed obligations to do.

3.2.2 The Generality of *D*-Fittingness

I have said that our *d*-reasons are derived from our desires to be *d*-fitting or avoid *d*-sanction-liability. This raises the question of the *individuation conditions* of the concept of *d*-fittingness. Even if the property of *d*-fittingness for one system of moral norms comes apart from the property of *d*-fittingness for another system of moral norms, we can still ask whether the concept of *d*-fittingness which fills in each of our desires are distinct. To see what I have in mind, consider

³²Well, perhaps. This depends on whether it being pointed out to you that you're in violation of what you ought_{*d*} to do is itself a sanction. I put this possibility aside for my purposes here.

³³Unless we are in a context in which the hands of the judges and police are tied. See (Woods 2016 §2).

an evolving system of legal norms. Even if we think the property of legal-fittingness evolves over the course of the change of the legal norms (I make no commitment here), we want it to be the case that the concept of legal-fittingness and the corresponding notion of law-abiding are stable over such minor evolution, even if the facts about whether we are actually behaving in a legally-fitting manner change. Presumably, if we make a radical change in our legal norms, we would get a corresponding difference in our concept of legal-fittingness, but the change would have to be pretty radical.

We could possibly treat notions like legal-fittingness in terms of “coherence with the salient normative system in the environment”, but I think this would be a mistake. This is not how notions of fittingness feel from the inside and it presumes that we recognize that normative systems vary or could from context to context. While this is obviously true for some systems of norms, like legal norms, it is manifestly not the case for systems of norms like morality, epistemic norms, and the like. We typically conceive of substantive systems of norms as universal (§2.1.3) and this suggests that we do not want our notion of *d*-fittingness to change dramatically from day to day. We could avoid this problem by focusing on sanction-liability, since this concept is more often fixed, but even then this also seems like it could change from context to context.

Instead, it seems to me that the right solution is to treat the concept of moral fittingness or moral blame-liability as coarse-grained, just as we should treat legal fittingness and legal punishment-liability as coarse grained. This coarse-graining means that we can (a) desire to be morally upstanding without knowing exactly what this entails (this is quite often the case) and (b) roughly correctly treat the property of moral fittingness as uniform across contexts, even though in fact the properties are extensionally different. This seems even more the right solution when we look at the contours of many systems of norms which tend to contain rather vague permissions, forbiddings, favorings and disfavorings. Even games often are tolerant of slight deviations in the rules and our desire to be chess-fitting doesn’t seem to change whether we are playing chess-with-*en-passant* or chess-without.³⁴

3.3 Filling the Gaps — The Psychology of Norm-Internalization

We are still absent sufficient materials to fully recapture the intuitive notion of a substantive set of norms. Even if it turns out that we nearly always have desires to be morally fitting, this only approximates the real definition of a substantive system of norms and the idea that substantive systems of norms play a role in practical deliberation. It doesn’t yet capture the sense in which substantive systems of norms are contentful and universal. And, of course, given the sort of account that I’ve given, I don’t believe that substantive systems of norms are either contentful or universal. So why do we think that they are? Perhaps surprisingly, there is empirical work that points towards an explanation of this.

Keupp, Behne, and Rakoczy (2013, 2016) have provided evidence that the contours of norm internalization differ between intuitively task-oriented and non-task-oriented norms. Very young children (roughly age 3) will *over-imitate* both, protesting actions that do not perfectly mimic the method they had been shown of achieving the task even in task-oriented behaviors. That is, they over-imitate in their own behavior and protest non-imitation. As they get older (around age 5), they tend to cease to over-imitate with obviously task-oriented behaviors, but continue to protest deviations from non-obviously-task-oriented behaviors.

³⁴This could be developed in more technical detail, but the point should be clear enough.

What do I mean by a ‘obviously task-oriented behavior’? Keupp, Behne, and Rakoczy use the method of instruction to delineate this difference. An obviously task-oriented behavior is, for the purpose of their experiments, one where the child is instructed in how to do it with a focus on the instrumental end of the behavior. A non-task-oriented behavior is one where the *method* is the focus of the experiment. They write:

. In particular, the more the focus of an action demonstration is shifted toward the method of performing the action rather than on the action’s instrumental effects, the more a causally irrelevant Action Element A will be perceived as a conventionally essential and obligatory part of a bigger Action C—and, consequently, the more children normatively expect others to reproduce this part (“This is the way this activity is performed correctly; it is part of the overarching goal both to produce the effect and to produce it in this way”) even if its causal irrelevance is obvious (for a related line of argument regarding rational imitation, see Király, Csibra, and Gergely, 2013). (Keupp, Behne, and Rakoczy 2013, pg. 396)

What their results thus show is that when the slightly older child is in a position to suss out that the purpose of the behavior is to bring about an end, then they allow that there are variable ways of accomplishing this end. They do not criticize these ways of bringing about the end, though they continue to criticize people who violate the method they were shown for the non-obviously task-oriented behaviors. Younger children are incapable of distinguishing these two cases well and overimitate, criticizing in both cases.

As I said in my 2016, if this is correct, then we should expect that as we get better at identifying instrumental reasoning, such as that involved in task-oriented behaviors, we cease over-imitating, but maintain a concern with acting correctly in tasks which are method-directed, such as morality, etiquette, and formation of beliefs. Moral norms and epistemic norms do not present themselves as achieving a goal or task, so presumably we tend not to treat them as task-oriented. This would mean that we tend to treat them as the sorts of thing to be matched nearly perfectly, not the sort of thing that we could achieve in many different ways.

This sits quite nicely with our view of substantive norms. Etiquette, unlike morality, is typically seen as a merely formal system of norms. I’ve always thought this undersells the importance of etiquette—in many contexts and in the past, etiquette has had a seemingly substantive status. Consider, for example, the Victorian England of myth and legend or honor cultures more generally. However, it seems to me explicable given the empirical hypothesis. We are trained up to think of etiquette as a task-oriented behavior. Etiquette is an attempt to coordinate on a social problem of reasonable interaction with each other. If we aren’t told this explicitly, we will experience it as soon as we engage with other ways of solving these social problem by living or engaging with different cultures. This line of thought would predict that when we are not widely-traveled, well-educated or -read, then we will treat etiquette as a substantive system of norms and our etiquette obligations as substantive. I put it to the reader that folks do exactly this. When we do not tag a system of norms with a task, we treat it as substantive. When we have been acculturated to do so, we treat it as merely formal.

3.3.1 Recapturing the Intuitive Notion of Substantive Obligation

Why, then, does our intuitive conception of a substantive system of norms involve contentfulness and universality? Exactly because we tend to view them as not being task-oriented behaviors and internalize them in a similar way to how we early on internalize even task-oriented behav-

iors/ We treat them as the sort of thing that we need to comply more or less exactly with, on pain of violating being how we ought simpliciter to act. We will treat them as contentful and universal since there is no obvious task which could be satisfied differently from context to context. To put this clearly, since we typically have non-instrumental desire to be morally fitting or to avoid moral-sanction-liability (which will depend on various further features of our psychology and environment), our moral obligations really are substantive in my above sense. But we treat morality as substantive exactly because we have been trained up to care about the *method* of being moral, so we do not think of moral fittingness as the sort of thing we could satisfy in various different ways.

Note that when there is a obvious variably-satisfiable task inherent in moral norms, such as with impure moral norms like “don’t call people names”, we are plausibly far more likely to treat these like we treat merely formal obligations. The task of this impure norm is to avoid hurt and, moreover, we eventually recognize this background task. When we find that we would not hurt people by calling them names—such as recognizing our culture is one where people regularly slag each other off—then we think that we need additional reason to comply with it. This sits nicely with the fact that some of us tend to become more morally permissive as we get older. When we see that the fringes of morality have some purpose to them and we see that differing impure moral norms are in force in different areas and seem to work fairly well, we tend to become less fetishistic about impure morality. The psychological results I’m leaning on here would suggest that what is going here is that we start to see impure morality as encoding an end to be achieved and, consequently, we start to see that that end could be achieved in various ways.

In short, our view of a portion of a system of norms as substantive seems to co-track with how explicit and obvious it is that the system of norms is directed at a particular end which itself will be explicable in terms our early normative education. This, of course, does not mean that we take these portions of systems of norms which we see as non-substantive to be the sort of thing that we can simply ignore when there’s an equally good way of getting to the same end. We still view legality, games, and etiquette as giving rise to obligations which we, typically, care about. And sanction-liability does not lapse when we view a system of norms as not substantive, as I have argued above.

If this is right, then it is no surprise that there is a tendency to treat substantive norms as both contentful and universal: if our *d*-concepts aren’t obviously beholden to a task, then there is no task for *d* that we take to be variably satisfiable. There is then there is no structure to change from context to context, and consequently, by default no difference a change in context will make to our view of whether or not a system of norms is in play. So we can recapture the intuitive notion of a substantive system of norms as a spandrel of early norm internalization where we tend to treat norms as universally binding *unless* we come to recognize them as a task-oriented behavior. No task, no structure to change from context to context. No structure to change from context to context, no temptation for us to believe that they are merely local norms. So, by default, we treat them as universally binding and as describing real facts about how we ought to behave morally or epistemically.

We are aided in this fact by the etiological explanation of the coarse-grained nature of the property of moral-fittingness. The moral instruction many of us received—our norm internalization behavior—was of the “don’t do this, it’s immoral” or “only (morally) bad people do such and so” variety. In particular, these labels were applied not only to those in our community, but

also to those abroad and in evaluating the historical past and envisioned future. It is thus no surprise that we tend to treat the concept of being morally upstanding or gross moral turpitude as universal contentful properties applying across contexts and that we came to care that we ourselves were morally upstanding. We internalized these concepts as “the way to do things” and, given the rough uniformity of moral behavior across contexts, we did not relinquish this internalization as we grew more sophisticated and experienced other cultures. We never had any explanatory need to seek out a purpose for the pure portion of our systems of norms. Whereas our universal and coarse-grained concept of properly polite behavior exploded once we experienced the dramatic distinctions in politeness norms and came to need to recognize how we could easily and variably we could accomplish the purpose that etiquette norms served to advance.

4 Conclusions and Coda

I have now suggested a view of the distinction between merely formal and substantive systems of norms and the obligations grounded thereby. I want to close by suggesting a way in which my account can be generalized, though at some cost. I have used desire-based reasons to motivate my account and defined a substantive system of norms and a substantive obligation in terms of those systems of norms d where we typically have desires to have d -fitting behavior. This is a nice naturalistic account and one which seems to me naturally connected to a humane view which provides a beautiful explanation of why we have reason simpliciter to do those things we desire to do—at base, that is what it is to have a reason at all.

However, for those who find this kind of view uncomfortable, note that we could rejigger almost everything I’ve said here in terms of a distinct source of reason. Suppose, for illustrative purposes, we took value-based reasons of the type suggested by Barry Maguire (2016) and Wedgwood (2009) as our fundamental source of reasons simpliciter. We could then define a substantive system of norms as a system of norms d where we typically (or always) have intrinsic reason to have d -fitting behavior. We could likewise say that a merely formal system of norms d is one where we do not have intrinsic reason to have d -fitting behavior or to avoid d -sanction-liability. And, finally, we could define a substantivish system of norms d as one where we have intrinsic reason to avoid having d -sanction-liability behavior. My worry for this sort of account is that it obscures the sorts of distinction in the reasons we have; it would, after all, be very bizarre for there to be *intrinsic* reason, in terms of their being final value, in not being d -sanction-liable. Whereas it is completely plausible that we could desire this intrinsically and, thereby, have reason to act as we ought $_d$.

A hybrid view (inspired by Maguire and Woods (ms)) which I think more plausible holds that a substantive system of norms d is one where it is finally valuable to desire to have d -fitting behavior. We can then run our analysis one level down and make use of the flexibility of desires in capturing the differences between our reasons to have d -fitting behavior and to avoid d -sanction-liability.³⁵ I can also offer my account of the difference between substantive and merely formal systems of norms to those who think there are non-desire-based reasons; simply use those desires we have reason simpliciter to have to characterize substantive obligation.

³⁵I have worries about this account, but they are more worries about the nature of value than worries about the account’s structural integrity.

References

- Chrisman, M. 2015. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press.
- Copp, D. 2004. Moral Naturalism and Three Grades of Normativity. *Normativity and naturalism*: 7–45.
- Finlay, S. 2007. Responding to Normativity. In *Oxford Studies in Metaethics, Vol. 2*, ed. R. Shafer-Landau, 220–39.
- Finlay, S. 2014. *Confusion of Tongues: A Theory of Normative Language*. OUP.
- Foot, P. 1972. Morality as a System of Hypothetical Imperatives. *The Philosophical Review*: 305–316.
- Grice, H. P. 1989. *Studies in the Way of Words*. Harvard University Press.
- Harman, E. 2016. Morally Permissible Moral Mistakes. *Ethics* 126(2): 366–393.
- Hart, H. L. A. 1960. Prolegomenon to the Principles of Punishment. *Proceedings of the Aristotelian Society* 60(1): 1–1.
- Hart, H. L. A. 1961. *The Concept of Law*. Oxford, Clarendon Press.
- Joyce, R. 2001. *The Myth of Morality*. Cambridge University Press.
- Keupp, S., C. Bancken, J. Schillmöller, H. Rakoczy, and T. Behne. 2016. Rational Over-Imitation: Preschoolers Consider Material Costs and Copy Causally Irrelevant Actions Selectively. *Cognition* 147: 85–92.
- Keupp, S., T. Behne, and H. Rakoczy. 2013. Why Do Children Overimitate? Normativity is Crucial. *Journal of Experimental Child Psychology* 116(2): 392–406.
- Kolodny, N. 2005. Why Be Rational? *Mind* 114(455): 509–563.
- Korsgaard, C. M. 1997. The Normativity of Instrumental Reason. In *Ethics and Practical Reason*, eds. G. Cullity and B. Gaut. Oxford: Clarendon Press.
- Kratzer, A. 1977. What 'Must' and 'Can' Must and Can Mean. *Linguistics and Philosophy* 1(3): 337–355.
- Maguire, B. 2016. The Value-Based Theory of Reasons. *Ergo*.
- Maguire, B. and J. Woods. manuscript. Explaining Epistemic Normativity.
- McPherson, T. 2011. Against Quietist Normative Realism. *Philosophical Studies* 154(2): 223–240.
- Parfit, D. 2011. *On What Matters*. Oxford University Press.
- Railton, P. 1984. Alienation, Consequentialism, and the Demands of Morality. *Philosophy and Public Affairs* 13(2): 134–171.
- Rawls, J. 1955. Two Concepts of Rules. *Philosophical Review* 64(1): 3–32.
- Schafer, K. 2016. The Modesty of the Moral Point of View. In *Weighing Reasons*. Oxford University Press.
- Shapiro, S. 2011. *Legality*. Harvard University Press.
- Sinclair, N. 2016. On the Connection between Normative Reasons and the Possibility of Acting for those Reasons. *Ethical Theory and Moral Practice*.
- Tiffany, E. 2007. Deflationary Normative Pluralism. *Canadian Journal of Philosophy* 37(5): 231–262.

Wedgwood, R. 2009. Intrinsic Values and Reasons for Action. In *Metaethics*, eds. E. Sosa and E. Villanueva, 342–363. Wiley Periodicals, Inc.

Woods, J. forthcoming. The Normative Force of Promising. *Oxford Studies in Normative Ethics*, vol. 6.