

This is a pre-print. The final version will appear in *Philosophy and Public Affairs*.

DOES “OUGHT” IMPLY “FEASIBLE?”*

Nicholas Southwood

I. INTRODUCTION

Imagine a mercantile and materialistic polity, Pecunia, in which achieving a modest improvement in the economic position of the poor by progressive taxation is infeasible inasmuch as most members of the middle-class majority simply cannot bring themselves to work at the requisite level of intensity for less pecuniary gain (cf Cohen 2008, ch. 1). Now consider the claim that

- (1) The Pecunians ought to improve the economic position of the poor by progressive taxation.

What should our reaction be to normative claims such as (1) that make infeasible demands?

The reaction that many of us *in fact* have, I take it, is to feel deeply conflicted (see Nagel 1991, ch. 2). On the one hand, claims such as (1) seem to be evidently *false* in virtue of demanding the infeasible (Goodin and Pettit 1995; Brennan and Southwood 2007; Southwood and Wiens forthcoming; Gilibert 2012; Gilibert and Lawford-Smith 2012; Wiens 2015). While it would be nice to imagine a world where members of the middle-class were disposed to be less concerned with their own lot and more concerned with the lot of the less fortunate, this is not the world we live in. In the world such as it is, where most members of the middle-class cannot bring themselves to do what is required to improve the economic situation of the poor by progressive taxation, it is simply false that the Pecunians ought to improve the economic situation of the poor by progressive taxation. To insist otherwise would amount to objectionable *unworldliness* – to chasing “pies in the sky.”

On the other hand, claims such as (1) also seem to be evidently *true* in spite of making infeasible demands (Estlund 2011, pp. 219-21; Cohen 2008; Gheaus 2013). The case of the Pecunians seems to be aptly described as a case where objectionable character traits of the

* For invaluable comments on and discussion about previous versions of this article I would like to thank Robert Audi, Geoffrey Brennan, Garrett Cullity, Philippe Chuard, Ramon Das, William Dunaway, David Estlund, Justin Fisher, Pablo Gilibert, Robert Goodin, Elizabeth Harman, Robert Howell, Karen Jones, Simon Keller, Suzy Killmister, Alex King, Holly Lawford-Smith, Seth Lazar, Sarah McGrath, Tristram McPherson, Emily McTernan, Glen Pettigrove, David Plunkett, Jonathan Quong, Luke Robinson, Geoffrey Sayre-McCord, Philip Pettit, Paolo Santorio, David Schmidt, Michael Smith, Kai Spiekermann, Daniel Star, Lachlan Umbers, Laura Valentini, Jonathan Way, David Wiens, and audiences at the Varieties of Normativity Workshop at Princeton University, the Southern Methodist University, Boston University, the Australasian Workshop in Moral Philosophy, and the Reasons and Virtues Conference at the Australian Catholic University. I am especially grateful to two anonymous referees for their razor-sharp comments and criticisms. Research for the article was supported by DP120101507 and DP140102468.

middle-class majority make it the case that it is infeasible for the Pecunians to do what they plainly *ought* to do. Insofar as the Pecunians fail to improve the economic situation of the poor by progressive taxation, we should say that they have done something *wrong*. To deny this would be to be guilty of treating the Pecunians with undue *lenience* – of mistakenly letting them “off the hook.”

Not everyone will share my ambivalence. But suppose you do. What is going on? A natural response is that what is going on is that we are making a mistake. What I shall call the principle that “ought” implies “feasible” (OF) holds that

(OF) An agent or set of agents A ought to realize a state of affairs x only if it is feasible for A to realize x (Brennan and Southwood 2007; Southwood and Wiens forthcoming).

(OF) is either true or false. Either it’s true, in which case our sense that (1) is true in spite of demanding the infeasible – and that denying the truth of (1) involves treating the Pecunians with undue lenience – involves a mistake. Or it’s false, in which case it is our sense that (1) is false on account of demanding the infeasible – and that insisting on the truth of (1) involves objectionable unworldliness – that involves a mistake. Either way, our ambivalence involves a mistake. The philosophical task is to say what the mistake is and why we are disposed to make it.¹

My aim in this article is to develop a different response. This begins from the supposition that our normative talk and thought encompasses a *plurality of oughts* – in particular, a plurality of oughts that are supposed to be capable of operating in the service of, and playing distinct roles associated with, what I shall call different *core normative practices* (cf Jackson 1991, pp. 471-72; Parfit 2011, pp. 35-6; Schroeder 2011).² The claim that our ambivalence must involve a mistake is tempting insofar as we are assuming that there is some *single* privileged ought at play in the idea that “ought” implies “feasible.” But the supposition that there is a *plurality* of oughts that are supposed to be capable of playing importantly different roles naturally points toward an alternative. The alternative holds

¹ One possibility is that we are mistakenly conflating the true principle that “ought” implies “feasible” with some related but false idea: e.g. that it’s not the case that we ought to *try* to do what is infeasible (Southwood and Wiens forthcoming, section 5). Another possibility is that we are mistakenly conflating the false principle that “ought” implies “feasible” with some true related but distinct idea: either a) that “ought” is related to “feasible” in some other way (e.g. that feasibility is (sometimes) *substantively significant* (Brennan and Pettit 2005; Cohen 2008; Estlund 2011; Lawford-Smith 2013); or that “ought” *conversationally implies* “feasible” (c.f. Sinnott-Armstrong 1984); or that infeasibility is *excusing*); or b) that “ought” implies (or is otherwise related to) something *like* “feasible,” such as “not too demanding” (cf Nagel 1991, ch. 3) or “not too difficult.” I confess that I don’t find this response – in either form – remotely plausible. The problem is that our ambivalent reactions simply seem more persistent and powerful than any such attempt to explain them away. But I won’t try to argue for that here.

² To anticipate, the core normative practices that I shall focus on are prescriptive, evaluative, deliberative and hypological practices.

that whether or not “ought” implies “feasible” depends on which core practice-serving ought is at issue. There is some salient core practice-serving ought for which it’s true that “ought” implies “feasible.” There is some other salient core practice-serving ought for which it’s false that “ought” implies “feasible.” The interpretation of those normative claims that seems false on account of demanding the infeasible – and for which it seems objectionably unworldly to insist otherwise – involves the former. The interpretation of the claims that seems true in spite of demanding the infeasible – and for which it seems unduly lenient to insist otherwise – involves the latter. Our ambivalence, far from involving any mistake, is entirely appropriate.

This is the possibility I wish to explore in detail.³ After offering a more precise characterization of the notion of feasibility that is at play (section II), I shall sketch a framework for understanding different core practice-serving oughts in general (section III) and then use this framework to consider which particular core practice-serving oughts, if any, might be capable of vindicating our ambivalence. I shall consider and reject a prevalent and prima facie promising account according to which the relevant oughts are the *prescriptive* ought and the *evaluative* ought (section IV). I shall then propose a different account that holds that the oughts we need are instead the *deliberative* ought (section V) and the *hypological* ought (section VI).⁴ An account of this kind affords us a compelling way of vindicating our ambivalence. Or so I shall argue.

II. FEASIBILITY AS WHAT WE CAN BRING OURSELVES TO DO

There are many things we might mean by describing the realization of a particular state of affairs as “feasible” or “infeasible.”⁵ By “feasible” let me *stipulate* that I have in mind specifically what we *can bring ourselves to do* in the sense helpfully articulated by David Estlund in an important recent paper (Estlund 2011). For the sake of argument I shall assume, following Estlund, that it doesn’t mean simply what we *can do* in the sense of what we are *able* to do (Estlund 2011, pp. 219-21; see also Vihvelin 2004, pp. 443-4). It may be

³ A number of other theorists have also made suggestions along similar lines (see below, section IV). However, first, the accounts to date have been brief and gestural. In particular, virtually nothing has been said about how to understand the *nature* of the different oughts. But, without such an account it is impossible to evaluate the claim that particular oughts are such that it is true or false that “ought” implies “feasible.” Second, as I shall argue, the particular oughts that have been adduced (the prescriptive ought and the evaluative ought) are the wrong oughts to vindicate our ambivalence. My aim is to do better on both fronts.

⁴ I am not the first theorist to make use of this pair of oughts. On the contrary, they are sometimes evoked, at least implicitly, to make sense of the ambivalent reactions many of us have in cases such as the case of Professor Procrastinate (see Jackson and Pargetter 1986). But they have not been applied hitherto to the feasibility issue in particular.

⁵ There has been an enormous recent surge of interest in the notion of feasibility and its normative relevance. See, for example, Raikka 1998; Cohen 2003; 2008; Mason 2004; Brennan and Pettit 2005; Brennan and Southwood 2007; Southwood 2015; 2016; Southwood and Wiens forthcoming; Estlund 2007, ch. 14; 2011; 2014; Swift 2008; Gilabert 2009; 2011; forthcoming; Gilabert and Lawford-Smith 2012; Lawford-Smith 2012; 2013; Hamlin and Stemplowska 2012; Miller 2013; Gheaus 2013; Wiens 2013; 2014; 2015.

that I *can* (i.e. am *able* to) save the life of a fellow surfer who is being attacked by a four-foot bronze-whaler, yet cannot *bring myself* to do so. It may be that my wife and I *can* (i.e. are *able* to) give up our beautiful newborn baby for adoption, but cannot *bring ourselves* to do so.⁶ I take the notion of what we can bring ourselves to do to be a relatively intuitive one. But can we do something to *explicate* the notion? Estlund does not offer such an explication. I shall offer one on his behalf.⁷

Consider the following example:

Suppose Bill pleads that he is not required to refrain from dumping because he is motivationally unable to bring himself to do it. ... He is simply deeply selfish and so cannot thoroughly will to comply. Dumping his trash by the road is easier than wrapping it properly and putting it by the curb or taking it to the dump. He wishes he had more willpower, and yet he doesn't have it. Refraining is something he could, in all other respects, easily do, except that he can't thoroughly will to do it (Estlund 2011, p. 220).

The example of Bill is meant to be an example of someone who *can do* something but who *cannot bring himself* to do it. According to Estlund, when we say that an agent "can" do something, we are saying that she is *able* to act in that way. Estlund endorses a simple version of the so-called "conditional" account of ability (see Moore 1912, chs 1, 6). In particular, he holds that:

(2) An agent A can (is able to) X if and only if were A to try and not give up, then A would be sufficiently likely to X (Estlund 2011, p. 212).

Given (2), it follows straightforwardly that Bill can refrain from dumping his rubbish. That's because if Bill were to try to refrain from dumping his rubbish and not give up, then he would be virtually certain to succeed. However, (2) faces familiar counterexamples involving so-called "finkish dispositions" (Vihvelin 2004). So I propose instead a "dispositional" analogue of (2), i.e.

⁶ How is this supposed to work in the context of groups? In the case of abilities, Estlund has suggested what he calls the "ability bridge principle." This holds that: "Meeting some standard is within a [group's] abilities only if any individual behaviors that would be necessary if the [group] is to meet that standard are within those individuals' abilities" (Estlund 2011, p. 236). In the same spirit, I suggest what we might call the "feasibility bridge principle." This holds that: It is feasible for a group to meet some standard only if any individual behaviors that would be necessary if the group is to meet that standard are ones that the individuals can bring themselves to perform. To be sure, there are certain important complexities concerning how to conceptualize feasibility in the context of groups that I am simply setting aside here. For discussion see Lawford-Smith 2012.

⁷ I am grateful to Jonathan Quong for persuading me of the need to do more than offer an intuitive gloss.

(3) An agent A can (is able to) X if and only if A is disposed, insofar as A tries to X and does not give up trying to X, to (be sufficiently likely to) X (cf Vihvelin 2004, pp. 437-40; Southwood and Gilabert ms).⁸

Like (2), (3) also implies that Bill can refrain from dumping his rubbish. That's because Bill is disposed to be virtually certain to succeed in refraining from dumping his rubbish insofar as he tries and does not give up. But, unlike (2), (3) avoids finkish counterexamples – at least given plausible and widespread assumptions about the nature of dispositions.⁹

What about the claim that Bill cannot bring himself to refrain from dumping his rubbish? Estlund doesn't offer any account of what it means to say that an agent "can bring herself" to do something, though he does speak repeatedly and suggestively of "motivational capacities" (Estlund 2011, e.g. pp. 207, 208, 210 n4, 211-12, 226). Kadri Vihvelin (2004, pp. 443-44), who also employs the locution, offers some further instructive remarks in defending her version of the dispositional account of ability from a classic counterexample (Lehrer 1968, p. 32; Wolf 1990, p. 99). The counterexample involves an agent who "suffers from some pathological condition which either impairs her capacity to make a rational decision concerning some particular kind of action (e.g., claustrophobia, e.g., a pianist overcome by stage-fright) or impairs her capacity to make *any* rational decision about what to do. But if she (somehow, miraculously) decided (or chose, formed the intention, etc.) to perform the relevant action, she would succeed" (Vihvelin (2004, p. 443). We can also suppose that she is disposed to do so. According to Vihvelin, we should respond by insisting that the agent is perfectly *able* to act in the relevant ways. "She knows how to walk; her legs aren't paralyzed. She knows how to play piano; her fingers aren't broken." What the agent "lacks [is] the *ability to bring herself to do those things*; she cannot *use her reasoning ability to bring it about that she intentionally does these things*. Due to her phobia (stage fright, panic, etc.), she is unable to choose or try to act *according to her own conception of what counts as a good reason for acting*" (Vihvelin 2004, p. 443; italics added).

If we put these ideas together, we get the following: An agent can bring herself to perform some act only if she has certain capacities that, while not required for her to be able to perform the act as such, are required for a certain kind of *rational-volitional route* to

⁸ (3), like (2), faces other objections as well (e.g. Austin 1956, p. 218). For discussion, see e.g. Maier 2014 and Vihvelin 2004. For the most part these objections are not relevant for our purposes here, so I shall set them aside. One objection that *is* relevant is the classic counterexample to its sufficiency offered by Lehrer (1968, p. 32) and Wolf (1990, p. 99) that is discussed in the following paragraph.

⁹ (3) will be susceptible to finkish counterexamples only insofar as we understand the disposition in (3) in terms of the subjunctive conditional in (2). But C. B. Martin (1994), David Lewis (1997) and others taught us that, in general, it is a very bad idea to analyze dispositions in terms of subjunctive conditionals – precisely because such analyses are susceptible to counterexamples involving finks. Once we understand the disposition in (3) in the right (fink-proof) way, then (3) obviously won't be vulnerable to such counterexamples.

performing the action. I take it that at least three sorts of capacities are required in addition to the ability to act.

First, the agent must have the capacity to *decide* to perform the act on the basis of believing that she ought to do so. Suppose that one is a longstanding victim of domestic violence. One recognizes that one ought to leave one's abusive partner and yet one's self-conception has been eroded to the point that one is incapable of deciding to do what one knows one ought to do. Suppose that, unlike many others in one's situation, one has access to sufficiently robust personal networks that would make it the case that one is robustly disposed to succeed in leaving one's partner insofar as one tries and does not give up. So one is *able* to leave one's partner. Yet one is incapable of taking, as it were, the volitional first step of deciding to do what one recognizes one ought to do. Every time one revisits the issue, one finds oneself either deciding to stay, or at least deferring the decision to leave, in spite of one's better judgment.¹⁰

Second, the agent must have the capacity to *try* to perform the act. Suppose that one's colleague is having a heart-attack right in the middle of a glass-bottomed bridge thousands of feet above the ground while one watches in dismay from the side (cf Estlund 2011, p. 230). To survive, she needs someone to come and give her urgent CPR. One is able to save her life in the sense that one is robustly disposed to succeed in saving her life insofar as one tries to save her life and doesn't give up. Moreover, one may have and indeed exercise the capacity to decide to save her life by venturing out onto the bridge on the basis of believing that one ought to do so. "I'm going to do it!" one says to oneself. But one has a paralyzing case of acrophobia in consequence of which one is incapable of trying. Trying in this context involves somehow "activating" one's decision. And one is incapable of activating one's decision. Every time one goes to take a step towards the bridge, one finds oneself blocked by one's own psychology.

Third, the agent must have the capacity *not to give up* (trying to perform the act). Suppose that one has accepted to referee a paper for a journal (cf Jackson and Pargetter 1986). One is perfectly *able* to complete the refereeing assignment inasmuch as one is robustly disposed to complete it insofar as one tries and doesn't give up. Let's suppose, moreover, that one is perfectly capable of deciding to complete the task and trying to do so.

¹⁰ Does being able to bring oneself to perform an act really require such a fine-grained capacity – as opposed to the capacity to decide to perform the act simpliciter? Consider Unreflective Unwin, who happens to have a peculiar loathing for normative talk and thought in consequence of which he is incapable of deciding to visit his elderly great-aunt on the *basis* of believing that he ought to, but who nonetheless is perfectly capable of *deciding* to visit his great-aunt – so long as his thought processes are unencumbered by beliefs about ought and obligation (Southwood and Eriksson 2011, p. 200; cf Arpaly 2000). Let's grant that there is a sense in which we might well describe Unwin as "able to bring himself to visit his great-aunt." However, it is not the sense I have in mind. The sense I have in mind is the sense in which – to quote Vihvelin, once again – we are able "to choose or try to act according to [our] own conception of what counts as a good reason for acting" (Vihvelin 2004, p. 443). For that the capacity to decide is insufficient. I am grateful to a referee for persuading me that we need the fine-grained capacity.

Indeed, one exercises these capacities every morning, making oneself a nice cup of tea and sitting down at one's desk with the paper before one. Yet, while one is capable of deciding and for that matter trying, one is incapable of avoiding giving up. One's propensity to procrastinate is too potent to be overcome. The prospect of doing the daily crossword, informing oneself about the upcoming Ashes tour, arranging photos of one's recent trip to Vanuatu – doing anything but the damned review – proves simply too alluring.

We might wonder what it is for an agent to have these capacities. I tentatively suggest the following:

- (4) An agent A has the capacity to rationally decide to X only if A is disposed, insofar as A believes she ought to X, to (be sufficiently likely to) decide to X (on the basis of believing that she ought to X).
- (5) An agent A has the capacity to try to X only if A is disposed, insofar as A decides to X, to (be sufficiently likely to) try to X.
- (6) An agent A has the capacity not to give up trying to X only if A is disposed, insofar as A tries to X, (to be sufficiently likely) not to give up trying to X.¹¹

This gives us a somewhat clearer idea of how it might be that Bill is able to refrain from dumping his rubbish, yet unable to *bring himself* to refrain from doing so. Here are three possibilities. He is not disposed to be sufficiently likely to decide to refrain from dumping insofar as he believes that he ought to refrain from dumping. Or he is not disposed to be sufficiently likely to try to refrain from dumping insofar as he decides to refrain from dumping. Or he is not disposed to be sufficiently likely not to give up trying to refrain from dumping his rubbish insofar as he tries to refrain from dumping his rubbish.

III. ON THE PLURALITY OF OUGHTS

I am supposing that there is a plurality of oughts – in particular, a plurality of oughts that are supposed to be capable of operating in the service of, and playing distinct roles associated with, different core normative practices. This is obviously a controversial view. I shan't try to defend it. Rather, I shall be *supposing* it for the sake of seeing whether it affords us a persuasive way of vindicating the ambivalent reactions that many of us have to normative claims that demand the infeasible. But I do need to make it more precise. How am I supposing that we understand these different core practice-serving ought claims in

¹¹ As with (3), there remain potential counterexamples to (4), (5) and (6) as they stand. I shall assume that whatever modifications suffice to make (3) immune to these counterexamples should also be made to (4), (5) and (6).

general? And how am I supposing that we go about establishing whether particular core practice-serving ought claims do or don't imply "feasible?"¹²

The framework within which I shall be operating involves two key theses. The first thesis holds that ought claims involve different core practice-serving oughts in virtue of being *supposed to be fit to be used in different core normative practices*. Call this the *fitting use thesis*. The basic idea is that each of our core normative practices involves what I shall call a *canonical task* and a *canonical executor* of that task. For an ought claim to be fit to be used in a practice means to be fit to be used by a relevant canonical executor in the relevant canonical task associated with the practice. So, to say that an ought claim involves a particular core practice-serving ought is to say that the claim is supposed to be fit to be used – and, hence, that its truth depends on whether it is in fact fit to be used – by a relevant canonical executor in the relevant canonical task.

Take the so-called *prescriptive* ought. The core normative practice in the service of which claims involving the prescriptive ought are supposed to be capable of operating is our *prescriptive* practice. The canonical task associated with our prescriptive practice involves some special mode of *prescription*: say, giving advice, or making demands, or issuing injunctions. The canonical executor is an appropriately situated *prescriber*: for example, a would-be adviser who is privy to relevant information and who has the requisite "standing" to advise. So, the fitting use thesis says that an ought claim (of the form A ought to X) involves the prescriptive ought just in case a) it is supposed to be fit to be used by some such appropriately situated prescriber to prescribe that A performs X, and b) the claim is true only insofar as it in fact fit to be used in this way.

Or take the so-called *evaluative* ought. The core normative practice in the service of which claims involving the evaluative ought are supposed to be capable of operating is our *evaluative* practice. The canonical task associated with our evaluative practice is some special mode of *evaluation*: say, one that involves making claims about what would be good or desirable. The canonical executor is some appropriately situated *evaluator*: say, an evaluator who is privy to relevant information and impartial. So the fitting use thesis says that an ought claim (of the form A ought to X) involves the evaluative ought just in case a) it is supposed to be fit to be used by an appropriately situated evaluator to positively evaluate A's Xing, and b) the claim is true only insofar as it in fact fit to be thus used. And so on.

¹² What kind of semantic theory am I committing myself to in talking of "a plurality of oughts?" There are two possibilities, and I won't try to decide between them. The first possibility is that our talk and thought involving "ought" is literally ambiguous. There are different *concepts* of ought corresponding to the different practices in the service of which "ought" claims are supposed to be capable of operating (see Schroeder 2011). The second possibility is that "ought" claims can express different propositions without ambiguity. On this view, "ought" is univocal but relative to a "modal base" and an "ordering source" (Kratzer 1977), and the modal base of any ought claim includes an implicit specification of the practice in the service of which the claim is supposed to be capable of operating. Thanks to Billy Dunaway and a referee for help here.

Different ought claims differ in virtue of their supposed fitness to be used by relevant canonical executors in carrying out relevant canonical tasks.

Notice that the fitting use thesis helps to explain how ought judgments might be true insofar as they involve one kind of core practice-serving ought and false insofar as they involve some other kind of core practice-serving ought. Consider Mark Schroeder's case of Luckless Larry:

Luckless Larry ... has recently come by many misfortunes – his parents and siblings have recently passed away, his wife has divorced him to run off with a younger man, he has lost his job including his health insurance, and he has recently been diagnosed with kidney disease, which will require expensive treatment. Larry deserves to win the lottery, if anyone does. So if there is to be any justice in this world, then Larry ought to win the lottery (Schroeder 2011, p. 8).

Now consider the claim that

(7) Larry ought to win the lottery.

On the one hand, (7) might seem to be true. The fitting use thesis can help to explain why. For (7) does indeed seem to be fit to be used in evaluation. That is, it seems perfectly fit to be used by an appropriately situated evaluator to evaluate positively Larry's winning the lottery. This suggests that the interpretation of (7) that seems true is one that involves the evaluative ought. (7) seems true insofar as it is supposed to be fit to be used in evaluation. On the other hand, (7) might also seem to be false. Again, the fitting use thesis can explain why. For (7) seems unfit to be used in prescription. That is, it seems unfit to be used by an appropriately situated prescriber with the requisite standing to prescribe (e.g. to advise or demand or enjoin) Larry to win the lottery. This suggests that (7) is false insofar as it is supposed to be fit to be used in prescription – and, hence, false insofar as it involves the prescriptive ought.

It is important not to confuse the fitting use thesis with the quite different idea that ought claims involve some core practice-serving ought just in case they are *actually used* in the relevant practice. Call this the *actual use thesis*. Consider the claim that

(8) Malcolm Turnbull, the current Prime Minister of Australia, ought to lobby his conservative colleagues in favor of legalizing same-sex marriage in Australia.

Suppose that being a mere philosopher I will never have the opportunity to advise or recommend, still less demand or enjoin Turnbull to do anything – and, hence, that I will never have the opportunity to use (8) in prescribing. The actual use thesis implies that I will

never employ the prescriptive ought to judge that (8). The fitting use thesis does not have this implication. It is enough that my judgment involves a claim whose truth is dependent, at least in part, on whether it is *fit to be used* in advising, or demanding, or enjoining that Turnbull acts in this way, whether or not it is ever *actually* used in that way.

It is also important to distinguish the fitting use thesis from the idea that ought claims involve some core practice-serving ought just in case they are supposed to be fit to be used in the relevant practice by a *judge*. Call this *the judgment-relative fitting use thesis*. Consider Gary Watson's unsportsmanlike and vanquished squash-player who is virtually certain, insofar as she goes to shake the hand of her victorious opponent, to end up smashing him in the face (Watson 1982; Smith 1995). Suppose that Watson's squash-players judges that:

(9) Right now I ought to go away and cool off rather than go and shake my opponent's hand.

Suppose that prescription is essentially *second-personal* in the sense that it is only possible to advise or demand or enjoin *another agent* to perform some act; it is impossible to advise or demand or enjoin *oneself* to perform an act (contra Smith 1994, p. 151). In that case, the judgment-relative fitting use thesis would seem to suggest that it is impossible for Watson's squash-player to use the prescriptive ought in normative judgments such as (9). By contrast, the fitting use thesis does not have this implication. It is enough that her judgment involves a claim whose truth is dependent, at least in part, on whether it is fit to be used *by a relevant canonical executor*. In the case of claims that involve the prescriptive ought, the relevant canonical executor is an appropriately situated *prescriber* and, hence, some *other* agent.

What does it mean for an ought claim to be "fit to be used" in a practice?¹³ Consider an analogy (cf Aristotle 1970, book 2). Some materials (e.g. tiles, corrugated iron, thatch) are fit to be used in the practice of constructing the roof of a house. Others (e.g. Styrofoam, cardboard, calcium) are unfit to be used in this practice. What does it mean to say that a particular material is fit (or unfit) to be used in the practice of constructing the roof of a house? The obvious answer is that it means that the material has the right kind of physical *profile* to be used to achieve the *constitutive aim* of constructing the roof of a house. The constitutive aim of a practice is an aim such that the practice only counts as the particular practice it does in virtue of bona fide participants in the practice typically possessing and acting in pursuit of the aim. The constitutive aim of the practice of constructing the roof of a house is, roughly, to offer sustained protection from the elements. So, the materials that are fit to be used in constructing the roof of a house are fit to be thus used inasmuch as they

¹³ One approach would be to take the notion as a primitive and trust that you know what I'm talking about. (A number of philosophers who have employed the notion of "fitness" or "fittingness" have taken this approach (see e.g. McHugh and Way forthcoming; Chappell 2012).) I shall try to do better.

possess properties (durability, non-porousness, non-reactivity) that are required to achieve the aim of offering sustained protection from the elements. The materials that are unfit to be used in constructing the roof of a house are unfit inasmuch as they lack these properties.

I suggest that we understand the notion of “fitness to be used in a core normative practice” in much the same way. Just as the practice of constructing the roof of a house has a constitutive aim, core normative practices also have different constitutive aims (see Korsgaard 1996; 2009; Velleman 2000; Katsafanas 2013). And just as the fitness of a material to be used in the practice of constructing the roof of a house is a matter of whether it has the right kind of *physical* profile to be used to achieve the particular constitutive aim of constructing the roof of a house, the fitness of an ought claim to be used in a core normative practice is a matter of whether it has the right kind of *agential* profile to be used by a relevant canonical executor in the relevant canonical task to achieve the particular constitutive aim of the practice in question. This is the second key thesis of our framework. Call it the *constitutive aim thesis*.

The constitutive aim thesis offers a nice explanation of our verdicts about the fitness of the claim that Luckless Larry ought to win the lottery to be used in prescription and evaluation. The constitutive aim of prescription is presumably *practical*: to get the agent to perform the relevant act by heeding the prescription – for example, by following advice, or fulfilling a demand, or complying with an injunction. Our prescriptive practice only exists in the particular form it does because bona fide participants in the practice typically possess and act in pursuit of this aim.¹⁴ But the claim that Larry ought to win the lottery does not have the right kind of profile to be used to achieve this aim. Even if Larry wins the lottery, there is no prospect of his winning the lottery by following advice to win the lottery. We can only follow advice if we have certain practical capacities. Larry clearly does not possess the practical capacities that would be required to follow advice to win the lottery. By contrast, the constitutive aim of evaluation is *not* practical. Rather, it is presumably purely “*classificatory*,” as Thomas Nagel once nicely put it (Nagel 1971, p. 109). Thus, so long as Larry’s winning the lottery constitutes the right kind of subject-matter for relevant evaluative categories to be applicable to it, the claim that he ought to win the lottery may be perfectly fit to be used to positively evaluate his winning the lottery. The fact that he lacks certain practical capacities is neither here nor there.

If we put the fitting use thesis and the constitutive aim thesis together, this means we have a way of characterizing different core practice-serving ought claims. Such claims differ

¹⁴ This is, of course, perfectly consistent with cases of prescription where the prescriber happens to lack the aim of getting the prescribee to perform the relevant act by heeding the prescription. A referee offered the nice example of a lawyer who regards herself as being under a professional obligation to offer the correct advice to a highly counter-suggestible client, even though she knows that the best way to get her client to do what he ought is to advise him to do the opposite. When she acts in accordance with her obligation, she is nonetheless giving advice.

in virtue of their ostensive fitness to be used in different core normative practices, where fitness to be used in a normative practice is a matter of having the right kind of profile to be used to achieve the constitutive aim of the practice. It also means that we have a recipe for determining whether particular core practice-serving ought claims, thus characterized, do or don't imply "feasible." The key is to identify the constitutive aim of the practice and the profile that ought claims must possess in order to be usable in the practice to achieve the practice's constitutive aim. A particular class of core-practice serving ought claims will be such that "ought" implies "feasible" just in case being usable to achieve the practice's constitutive aim requires a profile that includes being constrained by what is feasible.

IV. THE PRESCRIPTIVE OUGHT AND THE EVALUATIVE OUGHT

We are looking for some pair of core practice-serving oughts that can vindicate our ambivalence in the face of normative claims that demand the infeasible. A natural thought is that the oughts we need are the prescriptive ought and the evaluative ought that we encountered above.¹⁵ As Pablo Gilabert puts it, claims involving the prescriptive ought and claims involving the evaluative ought

have a different relation to feasibility. Prescriptive claims are, and evaluative claims need not be, false when the obligations they mention cannot [feasibly] be fulfilled. ... The ... principle that "ought" implies ["feasible"] can be used to challenge prescriptive claims: since they take feasibility constraints as a necessary condition for what they demand, they may not demand what [is infeasible]. No such challenge would be cogent with respect to evaluative claims (Gilabert 2011, p. 56).¹⁶

The account in question comprises two theses. The first thesis is that the ought that we need to capture our sense that normative claims are *false* in virtue of making infeasible demands, and that insisting otherwise amounts to objectionable *unworldliness*, is the *prescriptive* ought. As we have seen, ought claims involve the prescriptive ought just in case they are supposed to be fit to be used by an appropriately situated prescriber in *prescription*: say, in advising, or recommending, or proposing, or enjoining, or directing, or instructing, or commanding the agent to act accordingly. Let's suppose, moreover, that the constitutive aim of prescription is *practical*: to get the agent to perform the act by following

¹⁵ An account of this kind is very prevalent. It seems to be at least implicitly endorsed in some form or other by many political theorists, including Cohen 2008, p. 251; Estlund 2011, pp. 224-25; Gheaus 2013, p. 457; Gilabert and Lawford-Smith 2012, pp. 818-19; Goodin and Pettit 1995, p. 1; Hamlin and Stemplowska 2012, pp. 52-8; Mason 2004, pp. 255-9; Pettit 2012, p. 126; Swift 2008; Wiens 2014.

¹⁶ Since I am assuming an interpretation of "can" such that "can" and "feasible" are distinct, I have replaced "can" with "feasible" in the quotation to avoid misunderstanding. This is innocuous. As Gilabert makes clear elsewhere (Gilabert forthcoming), he holds that "can X" implies "can bring oneself to X."

advice, heeding a recommendation, complying with a directive, or whatever. An ought claim will only have the right kind of profile to be usable to achieve this practical aim insofar as the agent can bring herself to perform the prescribed act. So it follows that claims involving the prescriptive ought must be constrained by what agents can bring themselves to do and, hence (given our stipulation), by what it is feasible for them to do. We therefore have an explanation of why, on one salient interpretation, the claim that

(1) The Pecunians ought to improve the economic position of the poor by progressive taxation

is indeed *false* on account of demanding the infeasible. Since claims involving the prescriptive ought are false on account of demanding the infeasible, (1) is false insofar and because it involves the prescriptive ought. We also have an explanation of why insisting otherwise would amount to objectionable *unworldliness*. Insofar as (1) involves the prescriptive ought, it is supposed to be fit to be used by an appropriately situated prescriber in advising (or recommending, or proposing, or enjoining, or directing, or instructing, or commanding) that the Pecunians improve the economic position of the poor by progressive taxation. And, surely, it would be hopelessly unworldly to use (1) to advise (or recommend, or propose, or enjoin, or direct, or instruct, or command) the Pecunians to do what they cannot bring themselves to do.

The second thesis is that the ought we need to capture our sense that normative claims may be *true* in spite of making infeasible demands, that failing to act accordingly would be to do *wrong*, and that insisting otherwise amounts to treating agents unduly *leniently*, is the *evaluative* ought. As we have seen, evaluative ought claims are ought claims that are supposed to be fit to be used by an appropriately situated evaluator in *evaluating*. There are different views about the relevant mode of evaluation. One view holds that the relevant mode of evaluation involves special claims about what *ought to be the case*, namely claims to the effect that it ought to be the case that agents perform certain acts (Gheaus 2013, pp. 456-7; Brennan and Southwood 2007, pp. 12-13; Chuard and Southwood 2009, p. 602).¹⁷ A second possibility is that the relevant mode of evaluation involves special *conditional* (or *hypothetical*) ought claims, namely claims about what agents ought (e.g. in the prescriptive sense) to do if it is (or if it were) feasible for them to do it (Gilbert 2011, p. 56).¹⁸ A third possibility is that the relevant mode of evaluation involves special claims

¹⁷ We might worry that it invites circularity to understand claims involving the evaluative ought in terms of a mode of evaluation whose content itself includes an ought. I shall set aside this worry in what follows.

¹⁸ A referee has suggested that it seems misleading to describe conditional or hypothetical ought claims as involving a mode of *evaluation* and that it would be more accurate to say that Gilbert is trying to account for our ambivalence in the face of normative claims that make infeasible demands by distinguishing between *categorical*

about how agents *ought to be*, namely claims to the effect that agents ought to be such that they perform certain acts. Since the constitutive aim of evaluation is *classificatory*, rather than practical (see Nagel 1971, p. 109), there is no reason to think that, in order for ought claims to be usable in any of these ways to achieve the constitutive aim of evaluation, agents must be *able* to perform the relevant acts – still less that agents are able to *bring themselves* to do so – and, hence, no reason to suppose that claims involving the evaluative ought must be constrained by what is feasible. So we also have an explanation of why, on a different salient interpretation, the claim that

(1) The Pecunians ought to improve the economic position of the poor by progressive taxation

might indeed be true in spite of demanding the infeasible. Since claims involving the evaluative ought are not constrained by what agents can bring themselves to do, (1) might very well be true insofar as it involves the evaluative ought.

Have we found the account we need to vindicate our ambivalence? No. Despite its prevalence and prima facie appeal, we should reject the account. Neither thesis bears closer scrutiny. The problem with the first thesis is that claims involving the prescriptive ought may be *true* in spite of demanding the infeasible. Whether claims involving the prescriptive ought are false on account of demanding the infeasible turns on whether being feasible is among the conditions that must be met in order for an ought claim to have the right kind of profile to be usable to achieve the constitutive aim of prescription. But being feasible is *not* among these conditions. To see this, consider a modification of the case of Selfish Bill. Suppose that, while Bill is unable to bring himself to refrain from dumping his rubbish, he is perfectly able, as it were, to be *brought* to refrain from dumping his rubbish by heeding the prescription of an appropriately situated prescriber. While doomed to dumping if left to his own deliberative devices, he is not impervious to a judicious prescriptive intervention from the right quarters. This is a case where the claim that Bill ought to refrain from dumping is perfectly usable to achieve the aim of getting Bill to refrain from dumping his rubbish by heeding a prescription, despite the fact that he cannot bring himself to do so. What this means is that Bill's being able to bring himself to refrain from dumping his rubbish is *not* among the conditions that must be met in order for the claim that he ought to refrain from dumping his rubbish to be usable to achieve the constitutive aim of prescription. It's just not

and *conditional/hypothetical* oughts. While sympathetic to this line of objection, I shall continue to use Gilbert's own way of describing his position.

the case that prescriptive ought claims are necessarily false in virtue of demanding that agents do what they cannot bring themselves to do.¹⁹

What about the second thesis? This holds that the salient ought for which it's false that "ought" implies "feasible" is the evaluative ought. The problem here is that, while it's surely right that claims involving the evaluative ought may be true in spite of making infeasible demands, the evaluative ought is not *the right ought* to make sense of our reaction to normative claims that make infeasible demands such as the claim that the Pecunians ought to improve the economic position of the poor by progressive taxation. The interpretation of the claim that seems true and that we are trying to vindicate is one according to which the Pecunians would be doing something *wrong* insofar as they fail to improve the economic position of the poor by progressive taxation; and that denying this claim would amount to inappropriately letting them off the hook. Recourse to the evaluative ought can't explain this (cf Wiens 2014, p. 302).

Consider, first, the idea that the mode of evaluation associated with claims involving the evaluative ought involves claims that it ought to be the case that agents perform certain acts. At most, this allows us to say that, inasmuch as the Pecunians fail to improve the economic position of the poor by progressive taxation, things are not *as they ought to be*; that the world has *fallen short of some ideal*; that the book of history has been written in a way that is presumptively a *cause for regret*. All that's surely true, but it doesn't come close to capturing our sense that the Pecunians are doing something wrong.

Next, what about the Gilabert proposal that the mode of evaluation associated with claims involving the evaluative ought involves claims about what agents ought (say, in the prescriptive sense) to do if it is (or if it were) feasible for them to do it? This fails in the same way. At most, it allows us to say that the Pecunians ought (in the prescriptive sense) to improve the economic position of the poor by progressive taxation if it is (or if it were) feasible for them to do so. Since it's not feasible for the Pecunians to improve the economic position of the poor by progressive taxation, it turns out that they are not doing anything *wrong* in failing to do so. It might be said that they are doing something wrong in the evaluative sense. But this turns out to be nothing over above doing a conditional or hypothetical wrong (in the prescriptive sense).

Finally, what about the view that the mode of evaluation associated with claims involving the evaluative ought involves claims to the effect that agents ought to be such that

¹⁹ Doesn't being able to heed a prescription imply being able to bring oneself to act in the prescribed way? No. Whereas being able to bring oneself to act implies having the capacity to decide to perform the act on the basis of believing that one ought to perform the act, being able to heed a prescription to perform some act does *not* require having the capacity to decide to perform the act on the basis of believing that one ought to perform the act. It is enough to have the capacity to decide to perform the act on the basis of believing that the *prescriber believes* that one ought to do so. Recognizing that a suitably situated other agent believes that this is what one ought to do (and perhaps having this belief communicated via a prescriptive intervention) might have the necessary oomph to motivate one where one's own belief does not.

they perform certain acts? This is more promising. Inasmuch as the Pecunians are not such that they work at the same level of intensity for less pecuniary gain, we can say something more than that things are not as they ought to be. In addition, we can hold that the Pecunians have *gone wrong* in the sense that they have failed *to be as they ought to be*. Moreover, we are not forced to make this assessment conditional on whether it is feasible for the Pecunians to work at the same level of intensity for less pecuniary gain – indeed, assuming that “ought to be” doesn’t imply “feasible to be,” on whether it is feasible for the Pecunians to make it the case that they are such that they work at the same level of intensity for less pecuniary gain. Still, such an account cannot *fully* capture our adverse reaction. Let’s concede the Pecunians ought to be such that they work at the same level of intensity for less pecuniary gain. They ought to be otherwise than they are: less selfish, less materialistic, and so on. But does this really exhaust their normative failings? Does pointing out their failings of character really make the thought that there is no failing of conduct any less misguided? Surely not. Our assessment that the Pecunians are failing to do what they ought to do does not simply evaporate once we acknowledge that they are failing to be as they ought to be. To capture adequately our ambivalence in the face of normative claims that make infeasible demands, a different pair of oughts is required.

V. THE DELIBERATIVE OUGHT

What kind of ought might be capable of vindicating our sense that normative claims are necessarily false in virtue of making infeasible demands and that insisting otherwise amounts to objectionable unworldliness? The oughts we have encountered so far – the prescriptive ought and the evaluative ought – are not up to the task. I shall now argue that the ought we need is what I shall call the *deliberative ought*. Deliberative ought claims are ought claims that are supposed to be fit to be used in practices of (practical) *deliberation*.

To get a flavor of the deliberative ought, return again to the case of Luckless Larry. Suppose that Larry is contemplating his ghastly predicament and planning the rest of his unenviable life. Suppose that in the course of doing so he concludes that he ought to win the lottery. His judgment involves the deliberative ought in the sense I have in mind just in case it is supposed to be fit to be used in deliberation. Following R. Jay Wallace, let’s say that practical deliberation involves “resolving [or settling] the question of what one is to do” (see Wallace 2014, p. 1).²⁰ Moreover, let’s say that for Larry to use the claim that he ought to win

²⁰ I shall remain relatively neutral concerning exactly how to understand the nature of this question. All that I shall assume is that, whatever the right way to understand it, we settle it by making a *decision* (see also Wallace 2014, p. 1). So, for example, to engage in practical deliberation, it is not enough to engage in some purely doxastic activity that merely involves undertaking to form normative or non-normative beliefs, insofar as these beliefs fail to entail corresponding decisions. Rather, one must be engaged in a process that involves undertaking to make a decision. This leaves open a number of possible views. The view I myself like holds that the question of what we are to do is a special *sui generis* practical question that is distinct from both the question of what we ought to do and the question of what we will do (see Hieronymi 2009; 2011; Southwood 2016; forthcoming. But

the lottery in practical deliberation would be for him to treat the claim as *decisive* in settling the question of what he is to do: that is, to *decide* to win the lottery on the basis that this is what (or so he believes) he ought to do (see also Schroeder 2011). The claim that Larry ought to win the lottery seems plainly *unfit* to be used in this way and, hence, false insofar as it is supposed to be fit to be used in deliberation. Since claims about what we deliberatively ought to do are ought claims that are supposed to be fit to be used in deliberation – and whose truth depends on whether they are in fact fit to be used in deliberation – it follows that Larry’s judgment that he ought to win the lottery is false insofar as it involves the deliberative ought.

It is important to emphasize that I take the deliberative ought to be crucially distinct from the *prescriptive* ought. Whereas prescriptive ought claims are supposed to be fit to be used in practices of prescription, deliberative ought claims are supposed to be fit to be used in practices of deliberation. In contrast, a number of philosophers take the deliberative ought to *encompass* the prescriptive ought (e.g. Smith 1994, p. 151; Schroeder 2011; Kieseewetter forthcoming, ch. 1).²¹ Call this the *assimilationist view*. In my terminology, the assimilationist view holds that claims involving the deliberative ought are supposed to fit to be used, not only in deliberation, but also in *prescription*. There are two ways of interpreting the assimilationist view. One is *disjunctive*. The disjunctive interpretation holds that claims involving the deliberative ought are only true insofar as they are fit to play *at least one of these roles*. The other is *conjunctive*. The conjunctive interpretation holds that claims involving the deliberative ought are only true insofar as they are fit to play *both roles*.

Here is a quick argument against the assimilationist view – in either interpretation. First, while fitness to be used in deliberation and fitness to be used in prescription will typically overlap, they may also come apart. I will argue shortly that fitness to be used in deliberation requires being constrained by what is feasible. Since I argued in the previous section that fitness to be used in prescription does *not* require being constrained by what is feasible, it follows that an ought claim that demands the infeasible may be fit to be used in prescription despite being unfit to be used in deliberation. It is also plausible that they come apart in the other direction. That is, an ought claim may be fit to be used in deliberation but unfit to be used in prescription. (Suppose that while an agent has all the capacities required to use the claim in deliberation, she has some pathological incapacity to follow advice.) Second, if fitness to be used in deliberation and fitness to be used in prescription may come apart in both directions, then both the disjunctive and conjunctive interpretations of the

one might instead hold that it is simply the question of what one ought to do (Scanlon 1998) or the question of what one will (or what it is possible for one to) do (Harman 1975; Velleman 2000; Wallace 2001) – perhaps under a special practical “mode of presentation.” I am very grateful to a referee for persuading me of the merits of neutrality on this point.

²¹ Schroeder (2011) and Kieseewetter (forthcoming) also appear to take the deliberative ought to encompass the *hypological* ought. Again, I take them to be importantly distinct.

assimilationist view have implausible implications. Take an ought claim that is fit to be used in prescription but unfit to be used in deliberation. In that case, the disjunctive interpretation implies that the claim is true insofar as it involves the deliberative ought, even though it is not fit to be used in deliberation. That seems clearly wrong. Next, take an ought claim that is fit to be used in deliberation but unfit to be used in prescription. In that case, the conjunctive interpretation implies that the claim is false insofar as it involves the deliberative ought. Again, this seems clearly wrong. To avoid such implications we should keep the deliberative and prescriptive ought firmly apart, as I shall do in what follows.

Might the deliberative ought be the ought we are looking for to vindicate our sense that ought claims are necessarily false in virtue of demanding the infeasible? Given that 1) claims about what we deliberatively ought to do are ought claims that are supposed to be fit to be used in deliberation and 2) ought claims are fit to be used in deliberation just in case they are usable in deliberation to achieve the constitutive aim of deliberation, the crucial question for our purposes is: What is the constitutive aim of practical deliberation? In particular, must an ought claim be constrained by what is feasible to be usable in practical deliberation to achieve its constitutive aim?

I suggest that at least part of the constitutive aim of practical deliberation is *to determine how we will act by deciding how to act* (cf Katsafanas 2011, pp. 120-24).²² This has three elements. First, it has a *practical* element; it encompasses *action*. In engaging in practical deliberation – say, about whether to pat a little kangaroo who is eating grass in my front yard – part of my aim must be precisely to *determine I will act*: whether I will perform the act of patting the kangaroo, or whether I will refrain from performing that act. This does not mean that I will *in fact* act in accordance with my decision (insofar as I make a decision at all). I might decide to pat the kangaroo and then be thwarted by the world (the kangaroo runs away before I can pat him); or by another agent (who physically prevents me from patting him); or by my own laziness or incontinence. I might change my mind. (“On second thoughts, those claws look rather sharp.”) Nor does it mean that if I do succeed in acting in accordance with my decision, my action constitutes *part* of my deliberation, an idea that Aristotle is sometimes accused of having advocated. Nonetheless, I take it that my *aim* must be to determine how I will act. Suppose that my aim is simply to make a decision without aiming to act as I decide – say, because I have been offered a substantial prize to make a *decision* to pat the kangaroo even though patting him will result in something extremely unpleasant, such as getting badly bitten or clawed (see Kavka 1983). This may well be possible. But it isn’t practical deliberation.

²² I don’t say that this is *all* that the constitutive aim of practical deliberation involves. Plausibly, it also has some kind of *reason-responsive* aspect. However, I won’t say anything about what kind of reason-responsiveness (if any) is required here.

Second, the constitutive aim of practical deliberation has a *decisional* element; it encompasses making a decision (see Gibbard 2003; Southwood 2010, ch. 4). Clearly, it is possible to engage in practical deliberation without *in fact* making a decision. I might deliberate about whether to pat the kangaroo and fail to make a decision either way. (“It’s very cute. But those claws also look pretty nasty. Ah, I don’t know what to do.”) But it is not possible to engage in practical deliberation without having the *aim* of making a decision. Suppose that, like Lafcadio (André Gide’s dubious protagonist in *Les Caves du Vatican*), I have the aim of performing a purely *spontaneous* (or unmotivated) action and act in pursuit of that aim. Even if this is possible, it’s not practical deliberation. Or suppose that I am engaging in purely *speculative* reflection about whether I ought to pat the kangaroo, without any thought of making a decision. (Perhaps I have already decided to pat the kangaroo and am not reopening the issue.) Again, this does not count as practical deliberation in the sense I have in mind.

Third, the constitutive aim of practical deliberation has an *authorial* element; it encompasses a special *relation of authorship* between our decisions and our actions. In deliberating about what to do, we do not merely aim to determine how we will act and to decide how to act. Rather, we aim to determine how we will act *by* deciding how to act. Our aim, in other words, is for our decisions (whatever they are) to determine how we act (see Katsafanas 2011). Here, too, this relation will clearly not always obtain. Even when we succeed in making a decision (say, to pat the kangaroo) and doing as we decide (patting the kangaroo), we may not act on the *basis* of our decision but, rather, out of fear of appearing cowardly, or because someone with an authoritative bearing commands us to pat the kangaroo, or whatever. Still, at least an important part of the *aim* of practical deliberation is for our decisions to determine how we will act. Any activity in which this feature is lacking from our aim falls short of practical deliberation. Or so I shall assume.

Suppose that I am right that the constitutive aim of practical deliberation includes determining how we will act by deciding how to act. What kind of agential profile must an ought claim (to the effect that we ought to perform some act) possess to be usable in practical deliberation to achieve this constitutive aim? In particular, what kinds of *capacities* must we possess?

First, it must be a claim such that one has the *ability* to perform the act. Consider the claim that

(8) I ought to play the Scarbo movement from Ravel’s *Gaspard de la nuit* this evening.

Suppose that I am an incurably mediocre pianist who is unable to play Chopsticks, let alone a piece as difficult as *Gaspard de la nuit*. In this case (8) is clearly not usable to achieve the

aim of determining how I will act, still less the aim of determining how I will act by deciding how to act.

Second, the claim must be a claim such that one has the capacity to *decide* to perform the relevant act on the basis that this is what (or so one believes) one ought to do. Suppose that one is an impoverished but besotted parent who is confronted with the prospect of giving up a beloved child for adoption into a wealthy family. Now consider the claim that

(9) I ought to give up my child for adoption.

To use an ought claim in practical deliberation is to treat it as decisive in settling the question of what one is to do, that is, to decide to perform the relevant act on the basis that this is what (or so one believes) one ought to do. But suppose that one is simply incapable, even insofar as one believes that this is what one ought to do, of deciding to give up one's child for adoption on that basis. In that case, (9) is not usable in practical deliberation at all, still less usable to achieve the constitutive aim of practical deliberation.

Third, the claim must also be such that one has the further capacities required to perform the relevant act *by* deciding to perform it: i.e. the capacity to activate one's decision by trying; and the capacity not to give up trying. Suppose that my friend is drowning while an enormous shark cruises ominously between the shore and my drowning friend. Consider the claim that

(10) I ought to save my drowning friend.

Suppose that I am capable of *deciding* to save my friend on the basis that this is what (or so I believe) I ought to do, and even *able* to save her (were I to try and not give up I would be virtually certain to succeed). But suppose that I have a terrible case of selachophobia in consequence of which I am incapable of activating the decision by trying, or incapable of avoiding giving up trying. My selachophobia constitutes an insurmountable impediment, not to my deciding to save my friend or indeed to saving my friend but to saving my friend by deciding to save her. Once again, this means that (10) is not usable to achieve the constitutive aim of practical deliberation, namely to determine how I will act *by* deciding what I am to do.

So, in order for an ought claim (to the effect that one ought to perform some act) to be usable in practical deliberation to achieve the constitutive aim of practical deliberation, one must possess i) the ability to perform the act, ii) the capacity to decide to perform the act on the basis of believing that one ought to perform it, iii) the capacity to try to perform the act, and iv) the capacity not to give up trying to perform the act. But, of course, these *just are* the capacities that are required in order for one to be able to bring oneself to perform an act –

and hence, given our stipulation that feasibility is just a matter of what we can bring ourselves to do, that are required in order for it to be feasible for one to perform the act. So, being constrained by what is feasible is indeed part of the profile that ought claims must possess in order to be usable in practical deliberation to achieve its constitutive aim. Given that deliberative ought claims just are ought claims whose truth depends on their fitness to be used in practical deliberation, and the fitness of ought claims to be used in practical deliberation is a matter of whether they are usable in practical deliberation to achieve the constitutive aim of practical deliberation, it follows that deliberative ought claims are ought claims whose truth depends, at least in part, on whether they make feasible demands.

Let me briefly consider two important objections.²³ The first objection is that the idea that “deliberative ought” implies “feasible” entails an implausible account of practical deliberation. Suppose that Selfish Bill is deliberating about whether to refrain from dumping his rubbish. And suppose that he deliberates as follows: “I am not disposed to be sufficiently likely to decide to refrain from dumping my rubbish insofar as I believe that I ought to refrain from dumping my rubbish. It is feasible for me to refrain from dumping my rubbish only if I am disposed to be sufficiently likely to decide to refrain from dumping my rubbish insofar as I believe that I ought to refrain from dumping my rubbish. So, it’s not feasible for me to refrain from dumping my rubbish. Since claims involving the deliberative ought imply “feasible,” it is not the case that I deliberatively ought to refrain from dumping my rubbish.” If I’m right that the deliberative ought is such that “ought” implies “feasible,” then it seems to follow that Bill could indeed deliberate in this way. But it would not be appropriate for Bill to deliberate in this way. So it cannot be the case that claims involving the deliberative ought (as I have characterized it) are constrained by what is feasible.

In order to evaluate this objection it is important to distinguish two issues. One issue concerns whether it is possible for a deliberator such as Bill to discover the falsity of an ought claim simply by recognizing that he is not disposed to be sufficiently likely to decide to act in accordance with it insofar as he believes the claim? Let’s say that I am committed to saying that this is indeed possible.²⁴ Is this a problem? It would certainly be a problem if I were committed to saying that it is possible for Bill to discover that he *does nothing wrong* in dumping his rubbish. Fortunately, I am not committed to that. Indeed, the argument in the next section will be precisely that there is another salient ought – the hypological ought – for which there is no such sound argument by which Bill might arrive at the conclusion

²³ I am very grateful to a referee for raising these objections.

²⁴ Even this is not quite right. My endorsement of (4) was avowedly “tentative.” I am open to the possibility that we should replace (4) with something weaker like

(4’) An agent A has the capacity to rationally decide to X only if it is not the case that A is (robustly) disposed, insofar as A believes that she ought to X, not to (be sufficiently likely to) decide to X (on the basis of believing that she ought to X).

that it is not the case that he ought to refrain from dumping his rubbish. Rather, I am committed to the claim that it is possible for Bill to discover the falsity of the claim that he *deliberatively* ought to refrain from dumping his rubbish. This seems to be the right result. Deliberative ought claims are distinguished precisely by their ostensive fitness to be used in practical deliberation. If Bill discovers that the claim that he ought to refrain from dumping his rubbish lacks some important feature that is required in order to be fit to be used in practical deliberation, we should surely expect it to be possible for him to discover that it is not the case that he deliberative ought to refrain from dumping his rubbish.

The other issue concerns whether it is appropriate to use the principle that “deliberative ought” implies “feasible” *within practical deliberation*. I am not committed to saying that this is appropriate; I simply don’t take a stand on the issue. My aim has been to provide an account of the deliberative ought. According to this account, an ought claim involves the deliberative ought just in case it is supposed to be fit to be used in practical deliberation. I have given some necessary conditions for what it takes for an ought claim to be fit to be used in practical deliberation. The claim that Bill ought to refrain from dumping his rubbish violates one of these necessary conditions. But I do not purport to have given an account of what it takes for practical deliberation to be appropriate or inappropriate. In particular, I have said nothing at all about how (or indeed whether) it is appropriate for a deliberator’s beliefs about feasibility (or its constituents) to enter into practical deliberation. For all I have said, it may well be entirely *inappropriate* for Bill to deliberate by explicitly using the principle that “deliberative ought” implies “feasible.”

Wouldn’t there be something inconsistent or paradoxical about, on the one hand, insisting on the truth of this principle and, on the other hand, denying that it would be appropriate for a deliberator to employ it in practical deliberation? No more so than there is something inconsistent or paradoxical about a consequentialist denying that it would be appropriate to employ consequentialism in practical deliberation. Of course, it would be nice to have an explanation of *why* it might be inappropriate for a deliberator to employ the principle that “ought” implies “feasible” in practical deliberation – if indeed it *is* inappropriate. Such an explanation lies beyond the ambition of the current paper. But there are various possibilities. One possibility is the distinctively Kantian idea that, in order to be conducted under “the idea of freedom,” practical deliberation must be conducted in complete isolation from our beliefs about what is feasible (Kant 1998, ch. 3). Another (less drastic) possibility is that the role of feasibility assessments with regard to appropriate practical deliberation is restricted to a kind of *agenda-setting* role. Appropriately responding to our assessments about what is feasibility requires engaging (or continuing to engage) in practical deliberation concerning only those acts that we take to be feasible (or that we don’t take to be infeasible). But it does not permit us to arrive at normative beliefs on the basis of assessments about what is infeasible *within* deliberative decision-making.

Apart from these brief and speculative remarks, I am afraid that I must leave the question of exactly what role feasibility assessments should play within (or with regard to) practical deliberation for another occasion.

The second objection is that even if I am right that claims involving the deliberative ought are false on account of making infeasible demands, recourse to the deliberative ought cannot offer an explanation of our sense that normative claims are false on account of making infeasible demands *outside first-personal deliberative contexts*. It's one thing to say that if *Bill* judges that he (Bill) ought to refrain from dumping his rubbish, then his judgment is false insofar as it involves the deliberative ought. But, suppose, instead, that Bill's neighbor, *Sheila*, judges that Bill ought to refrain from dumping his rubbish. It might be thought that Sheila's judgment cannot involve the deliberative ought. Deliberation is essentially *first-personal* (see Owens 2011). Barring very special circumstances, we cannot deliberate on behalf of others. We can deliberate *as if* we are others. We can deliberate about how to *advise* others or what to *make* them do (by force or persuasion). But we cannot engage in the distinctively deliberative activity of settling the question of how others are to behave (Southwood 2016). This might be thought to suggest that it is simply impossible for third-personal normative judgments such as Sheila's to involve the deliberative ought. If this is right, then appealing to the deliberative ought cannot explain how such judgments are supposed to be false on account of making infeasible demands; hence, our account is crucially incomplete.

Fortunately, however, it's not right. Recall that core practice-serving ought claims are ought claims that are supposed to be fit to be used by a *relevant canonical executor* of the canonical task associated with the relevant core normative practice. The canonical executor of the canonical task associated with our deliberative practices is the *agent* whose conduct is the subject of deliberation. Since Sheila is not the agent of Bill's conduct, she cannot *herself* use the claim that Bill ought to refrain from dumping his rubbish in practical deliberation. Nor is it plausible to suppose that her judgment involves an ought claim that is supposed to be *fit* to be used by her (Sheila) in practical deliberation. What she can do, however, is make a normative judgment that involves an ought claim that is supposed to be fit to be used by the agent (i.e. Bill) in practical deliberation. That is enough in order for her judgment to involve the deliberative ought in the sense I have in mind.

Isn't there something odd about the idea that we may make normative judgments that involve ought claims that are supposed to be fit to be used by other agents? Not at all. We have already seen that this is true of first-personal *prescriptive* ought judgments, as when Watson's cantankerous squash-player judges that she ought not to go and shake her opponent's hand. Such judgments involve the prescriptive ought insofar as they are supposed to be fit to be used by an appropriately situated prescriber in prescribing. Assuming that prescription is essentially second-personal, this will have to be another agent.

I take it that there is no oddity here. The point is that our talk and thought involving the prescriptive ought often reflects our concern with whether certain conduct is *advisable* or *unadvisable* – a concern that encompasses our own conduct as well as that of others. Similarly, our talk and thought involving the deliberative ought reflects our concern with whether certain conduct is, as we might put it, deliberatively *choice-worthy* or *non-choice-worthy*. It seems clear that this concern extends to others' conduct as well as our own (see Jackson 1991, pp. 462-63).

So we appear to have found an ought for which it is true that “ought” implies “feasible.” Is it the right ought to vindicate the first part of our ambivalence in the face of normative claims that demand the infeasible? It seems to me that it is. Consider again the claim that

- (1) The Pecunians ought to improve the economic position of the poor by progressive taxation.

I suggest that the interpretation of (1) that seems straightforwardly false – and for which insisting otherwise seems objectionably unworldly – is precisely an interpretation of (1) that is supposed to be fit to be used in practical deliberation. Insofar as it involves the deliberative ought, to insist upon the truth of (1) would be to be committed to claiming that (1) is fit to be treated by the Pecunians as decisive in settling the question of whether to improve the economic position of the poor by progressive taxation. But this would be to insist that (1) is fit to be treated as decisive in settling the question of what they are to do in favor of doing what they cannot bring themselves to do – surely the very epitome of unworldliness.

VI. THE HYPOLOGICAL OUGHT

I have argued that the deliberative ought is the ought we need to explain the first part of our ambivalence in the face of normative claims that demand the infeasible. The remaining task is to identify the ought that can explain the second part of our ambivalence: our sense that such claims may also be *true* in spite of making infeasible demands, and that insisting otherwise amounts to treating agents unduly *leniently*, or mistakenly letting them *off the hook*.

My suggestion is that the ought we need for this purpose is the *hypological ought* (cf Zimmerman 2006). Claims involving the hypological ought are supposed to be capable of operating in the service of our *hypological* practices: our practices of directing *criticism* toward others and ourselves.²⁵ More precisely, I shall understand claims about what we

²⁵ Does this mean that the hypological ought is simply a particular kind of *evaluative* ought? While the labels don't especially matter, I believe this would be a mistake. That's because the various different core normative practices are defined, in part, by their constitutive aims. I have been assuming that the constitutive aim of our

hypologically ought to do as ought claims that are supposed to be fit to be used by an appropriately situated²⁶ agent in practices of *prospective criticism*: namely, practices of undertaking to criticize agents – assuming that they are suitably well-informed about relevant matters of fact²⁷ – insofar as they fail to act accordingly. To show that claims involving the hypological ought may be true in spite of making infeasible demands, we must find some plausible specification of the constitutive aim of prospective criticism such that being feasible is *not* among the conditions that must be met in order for ought claims to be usable to achieve this aim. What might this look like?

One possibility is that the constitutive aim of prospective criticism – like the constitutive aim of prescription and deliberation – is *practical*: to get the agent to perform the relevant act. It follows relatively straightforwardly that claims involving the hypological ought, thus construed, may be true in spite of making infeasible demands. For it is clearly not necessary that one be able to bring oneself to perform an act in order for an ought claim to be usable by an appropriately situated prospective criticizer to achieve this aim. What is necessary instead is that one be able to be *brought* to perform the act by the prospect of criticism. But while this way of elucidating the hypological ought explains how normative claims that demand the infeasible may be true, it cannot do so in an especially satisfying way. To see this, consider a further modification of the case of Selfish Bill. Suppose that Bill has sunk to such depths of selfishness and laziness that he cannot even be *brought* to refrain from dumping his rubbish by the prospect of criticism. I take it that there is no temptation to alter our verdict that there is an important sense in which it is nonetheless *true* that Bill ought to refrain from dumping his rubbish, and that denying this would amount to mistakenly letting Bill off the hook. But if the constitutive aim of prospective criticism is to get the agent to perform the relevant act, then we cannot explain this further verdict by evoking the hypological ought. For it follows that the claim that Bill ought to refrain from dumping his rubbish is *unfit* to be used in undertaking to criticize Bill insofar as fails to dump his rubbish and, hence, false insofar as it involves the hypological ought.

The upshot is that if we are to recruit the hypological ought to vindicate the second part of our ambivalence in a fully satisfying way, we had better find a specification of the constitutive aim of prospective criticism that does *not* encompass getting the agent to

evaluative practices is purely “classificatory” (Nagel 1971, p. 109). But it is not plausible to suppose that the constitutive aim of our hypological practices is purely classificatory. I am grateful to a referee for raising this issue.

²⁶ I won’t try to say exactly what it takes for a prospective criticizer to be “appropriately situated.” However, plausibly it involves being informed about relevant matters of fact and having the requisite “standing” to criticize.

²⁷ By this I mean that agents are sufficiently well-informed about matters of fact that bear upon the substantive profile of the act in question. While I will generally omit this cumbersome qualification in what follows for ease of exposition, it is necessary to avoid a serious objection based on cases of *blameless wrongdoing*. I am very grateful to a referee for raising this objection and forcing me to clarify my account of the hypological ought accordingly.

perform the relevant act. What could this be? Here is a natural thought. Suppose that the constitutive aim of prospective criticism is simply to hold one another (and ourselves) *accountable* for how we act. It is now widely accepted that criticism amounts to a kind of “intangible” sanctioning mechanism that operates alongside the kinds of tangible sanctions that are associated with violations of important parts of the law (see Brennan and Pettit 2004; Brennan et al 2013, ch. 10). At least for well-functioning agents, criticism is something we care about. Doling it out amounts to a way of “enforcing” or “policing” compliance with certain norms. Just as the prospect of tangible legal sanctions operates to make us accountable to one another as fellow citizens, the prospect of intangible hypothetical sanctions operates to make us accountable to one another (and ourselves) as fellow agents.

Which conditions must be met in order for an ought claim to be usable in prospective criticism to hold an agent accountable? One possible condition is that the agent must be capable of registering, acknowledging, and being appropriately responsive to the prospect of criticism (see McGeer and Pettit 2015).²⁸ Another possible condition is that *there was some time at which it was true* that the agent could have taken steps to make it the case that, now, she is able to bring herself to perform the relevant act. Where these conditions are not met, it might be thought that the agent is not fully *answerable* for her conduct in the way that is required for us to be able to hold her properly accountable for her failure to perform a relevant act. Perhaps these conditions are valid, perhaps not. I’m not going to take a stand either way.

It is enough for my purposes to note that it is surely *not* necessary that the agent be able to bring herself, or be brought (by prospective criticism, or prescription, or whatever), to perform the relevant act. Take Selfish Bill. Bill cannot bring himself, or be brought, to refrain from dumping his rubbish. Suppose, moreover, that a) Bill is perfectly capable of registering, acknowledging, and being appropriately responsive to prospective criticism insofar as he dumps his rubbish (say, by undertaking to compensate his neighbours and to enrol in anti-dumping classes) and b) Bill could easily have taken steps to make it the case that, now, he is able to bring himself, and be brought, to refrain from dumping his rubbish. Under these circumstances, should we say that the claim that Bill ought to refrain from dumping his rubbish is unfit to be used in prospective criticism because using it in that way would be at odds with the aim of holding him accountable? On the contrary, it seems clear that holding Bill accountable requires precisely *holding firm* and undertaking to criticise him regardless of his inability to bring himself, or be brought, to refrain from dumping his

²⁸ One important way of being appropriately responsive to the prospect of legitimate criticism is to avoid the prospect eventuating by undertaking to perform the relevant act. Other ways may include undertaking to acknowledge wrongdoing, to apologise, to have or express certain reactive attitudes, to mitigate the adverse effects of one’s conduct, to compensate those adversely affected, and taking steps to alter what we can bring ourselves to do in the future.

rubbish. While using the claim that he ought to refrain from dumping his rubbish to undertake to criticize him has no prospect of getting him to refrain from dumping his rubbish, anything less would amount to a failure to hold him to account.²⁹

If I am right that the constitutive aim of prospective criticism is to hold one another (and ourselves) accountable, then feasibility is not among the conditions that must be met in order for an ought claim to be usable in prospective criticism to achieve its constitutive aim and, hence, fit to be used in prospective criticism. Claims involving the hypological ought, thus construed, may be true in spite of demanding that agents perform acts they cannot bring themselves to perform. Moreover, claims involving the hypological claim may also be true in spite of demanding that agents perform acts they cannot be *brought* to perform (by prospective criticism, or prescription, or whatever).

So we have found an ought – the hypological ought – that does not seem to be constrained by what is feasible. But is the hypological ought the *right* ought to vindicate the second part of our ambivalence in the face of certain normative claims that make infeasible demands? Consider again the claim that

- (1) The Pecunians ought to improve the economic position of the poor by progressive taxation.

Can recourse to the hypological ought vindicate our sense that (1) is true in spite of demanding the infeasible, that the Pecunians would do wrong insofar as they fail to improve the economic position of the poor by progressive taxation, and that insisting otherwise would amount to mistakenly letting them off the hook?

First, recourse to the hypological ought can obviously explain why (1) might be true in spite of demanding the infeasible. Suppose we interpret the ought in (1) as involving the hypological ought. Given the account of the hypological ought that I have offered, this means that (1) is supposed to be fit to be used in prospective criticism. And (1) may be perfectly fit to be used in prospectively criticizing the Pecunians insofar as they fail to improve the economic position of the poor by progressive taxation in spite of the fact that they cannot bring themselves to do what is necessary to improve the economic position of the poor by progressive taxation. Second, such an interpretation of (1) also nicely captures our sense that there is something *wrong* with the Pecunians' conduct insofar as they fail to improve the economic position of the poor by progressive taxation, not merely something

²⁹ I am not, of course, denying that holding an agent to account often has an important *practical* dimension. Indeed, in what we might call *favorable cases* – where the agent *is* able to bring herself, or be brought, to perform the act – holding her to account may involve precisely aiming to get her to avoid the prospect eventuating by performing the act. In favorable cases, this may well be the *uniquely* appropriate way for the agent to be responsive.

regrettable, or something conditionally wrong, or something wrong with their character. Their conduct is wrong in the sense that it is criticizable. Third, there is a clear sense in which denying (1) on the grounds that it is infeasible for the Pecunians to improve the economic position of the poor by progressive taxation would be to treat them with undue lenience. Indeed, talk of “mistakenly letting agents off the hook” suggests precisely refraining from prospective criticism when prospective criticism is warranted, thereby mistakenly failing to hold them to account.

VII. CONCLUSION

Many of us feel internally conflicted in the face of normative claims that demand the infeasible. This might be because we are making a mistake. My aim in this article has been to explore in detail the alternative hypothesis that our ambivalent reactions are entirely consistent and appropriate. Rather than some *single* privileged ought such that “ought” implies “feasible” is either true or false, there are simply different core practice-serving oughts that our ambivalent reactions may be tracking. I have offered a way of understanding these different core practice-serving oughts in general, considered and rejected one appealing existing account concerning the particular core practice-serving oughts that are at play, and proposed a different account.

Let me close by saying something very briefly about three assumptions that have played a significant role in my argument. First, I have assumed a stipulative definition of feasibility such that by “feasible” we mean specifically what we can bring ourselves to do. Suppose you reject this assumption. Does this mean that what I have said is devoid of interest? No. Both the general framework for understanding different core practice-serving oughts and the accounts of the particular oughts (such as the deliberative ought and the hypological ought) that I have offered might very well be put to work to interrogate whether “ought” implies “feasible” in some other sense.

Second, I have assumed that, like me, you are ambivalent about normative claims that demand the infeasible. Suppose you are not ambivalent. Rather, such claims strike you either as straightforwardly false on account of making infeasible demands, or as straightforwardly true in spite of making infeasible demands. Does this mean that what I have said is irrelevant? No. It might be that you have been implicitly (or even explicitly) privileging an interpretation of the relevant normative claims in which the relevant ought is the deliberative ought or the hypological ought. In that case, I have provided you with a way of making sense of the interpretation that you take to be privileged. Moreover, there is, as it were, an *interpersonal analogue* of our problem of explaining ambivalence in the face of normative claims that demand the infeasible: namely, to make sense of *disagreement* among theorists who are non-ambivalent concerning what to say about normative claims that

demand the infeasible. The account I have presented here might be used to show that neither side need be making an error.

Third, I have assumed a particular framework that holds that there is a plurality of core practice-serving oughts and that they are to be understood in a certain way. This framework is by no means uncontroversial. I have done nothing to argue for it directly. Obviously, doing so would be a formidable task. Does this mean that the conclusion of the paper rests on a framework for which we have provided no support? No. The conclusion of the paper itself provides some support for the framework. The fact that the framework affords us a compelling vindicating explanation of our ambivalence in the face of normative claims that demand the infeasible is itself some evidence that it is true.

References

- Aristotle. 1970. *Physics*, trans. W. Charlton. Oxford: Clarendon Press.
- Arpaly, Nomy. 2000. "On Acting Rationally Against One's Best Judgment," *Ethics*, 110, 488-513.
- Austin, J. L. 1956. "Ifs and Cans," *Proceedings of The British Academy*, 42, 107-132.
- Blake, Michael. 2001. "Distributive Justice, State Coercion, and Autonomy," *Philosophy & Public Affairs*, 30, 257-96.
- Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. 2013. *Explaining Norms*. Oxford: Oxford University Press.
- Brennan, Geoffrey and Philip Pettit. 2004. *The Economy of Esteem*. Oxford: Oxford University Press.
- Brennan, Geoffrey and Philip Pettit. 2005. "The Feasibility Issue," *The Oxford Handbook of Contemporary Philosophy*, ed. F. Jackson and M. Smith. Oxford: Oxford University Press.
- Brennan, Geoffrey and Nicholas Southwood. 2007. "Feasibility in Action and Attitude," *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, ed. T. Rønnow-Rasmussen, B. Petersson, J. Jönsson and D. Egonsson, 1-25.
- Chuard, Philippe and Nicholas Southwood. 2009. "Epistemic norms without voluntary control," *Nous*, 599-632.
- Cohen, G. A. 2003. "Facts and Principles," *Philosophy and Public Affairs*, 31, 211-45.
- Cohen, G. A. 2008. *Rescuing Justice and Equality*. Cambridge, Mass.: Harvard University Press.
- Estlund, David. 2007. *Democratic Authority: a Philosophical Framework*. Princeton, NJ: Princeton University Press.
- Estlund, David. 2011. "Human Nature and the Limits (If Any) of Political Philosophy," *Philosophy & Public Affairs*, 39, 207-37.
- Estlund, David. 2014. "Utopophobia," *Philosophy & Public Affairs*, 42, 113-34.
- Gheaus, Anca. 2013. "The Feasibility Constraint on the Concept of Justice," *Philosophical Quarterly*, 63, 445-64.

- Gibbard, Alan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gilbert, Pablo. 2009. "The Feasibility of Basic Socioeconomic Human Rights: a Conceptual Exploration," *Philosophical Quarterly*, 59, 659-81.
- Gilbert, Pablo. 2011. "Feasibility and Socialism," *The Journal of Political Philosophy*, 19, 52-63.
- Gilbert, Pablo and Holly Lawford-Smith. 2012. "Political Feasibility: A Conceptual Exploration," *Political Studies*, 60, 809-25.
- Goodin, Robert E. and Philip Pettit. 1995. "Introduction," *A Companion to Contemporary Political Philosophy*. Oxford: Blackwell.
- Hamlin, Alan and Zofia Stemplowska. 2012. "Theory, Ideal Theory, and the Theory of Ideals," *Political Studies Review*, 10, 48-62.
- Hayek, Friedrich. 1935. *Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism*. London: Routledge.
- Hieronymi, Pamela. 2009. "The Will as Reason," *Philosophical Perspectives*, 23, 201-220.
- Hieronymi, Pamela. 2011. "Reasons for Action," *Proceedings of the Aristotelian Society*, 111, 407-27.
- Jackson, Frank. 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics*, 101, 461-482.
- Jackson, Frank and Robert Pargetter. 1986. "Oughts, Options, and Actualism," *The Philosophical Review*, 95, 233-55.
- Kamminga, Menno. 2006. "Why Global Distributive Justice Cannot Work," *Acta Politica*, 41, 21-40.
- Kant, Immanuel. 1998. *Groundwork of the Metaphysics of Morals*, trans. Mary Gregor. Cambridge, UK; New York: Cambridge University Press.
- Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics: Nietzschean Constitutivism*. New York: Oxford University Press.
- Kavka, Gregory. 1983. "The Toxin Puzzle," *Analysis*, 43, 33-6.
- Kiesewetter, Benjamin. Forthcoming. *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kolodny, Niko. 2005. "Why Be Rational?" *Mind*, 114, 509-63.
- Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Korsgaard, Christine. 2009. *Self Constitution: Agency, Identity, and Integrity*. Oxford; New York: Oxford University Press.
- Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean," *Linguistics and Philosophy*, 1, 337-355.

- Lawford-Smith, Holly. 2012. "The Feasibility of Collectives' Action," *The Australasian Journal of Philosophy*, 90, 453-67.
- Lawford-Smith, Holly. 2013. "Understanding Political Feasibility," *Journal of Political Philosophy*, 21, 243-59.
- Lehrer, Keith, 1968. "Cans without Ifs," *Analysis*, 29, 29-32.
- McGeer, Victoria and Philip Pettit. 2015. "The Hard Problem of Agency and Responsibility," *Oxford Studies in Agency and Responsibility*, 3, 160-88.
- McHugh, Conor and Jonathan Way. Forthcoming. "Fittingness First," *Ethics*.
- McPherson, Tristram. Ms. "The Truth in Constitutivism," unpublished paper.
- Maier, John. 2014. "Abilities," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, <<http://plato.stanford.edu/archives/fall2014/entries/abilities/>>.
- Mason, Andrew. 2004. "Just Constraints," *British Journal of Political Science*, 34, 251-68.
- Miller, David. 1976. *Social Justice*. Oxford: Oxford University Press.
- Miller, David. 2005. "Against Global Egalitarianism," *The Journal of Ethics*, 9, 55-79.
- Miller, David. 2013. *Justice for Earthlings: Essays in Political Philosophy*. Cambridge, UK: Cambridge University Press.
- Moore, G. E. 1912. *Ethics*. London: Williams and Norgate.
- Nagel, Thomas. 1971. *The Possibility of Altruism*. Oxford: Clarendon.
- Nagel, Thomas. 1991. *Equality and Partiality*. Oxford and New York: Oxford University Press.
- Nove, Alec. 1991. *The Economics of Feasible Socialism Revisited*. London: Harper Collins.
- Owens, David. 2011. 'Deliberation and the First Person'. In *Self Knowledge*, ed. Anthony Hatzimoysis. Oxford: Oxford University Press, pp. 261-77.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- Pettit, Philip. 2012. *On the People's Terms*. Cambridge: Cambridge University Press.
- Portmore, Douglas. 2012. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford; New York: Oxford University Press.
- Raikka, Juha. 1998. "The Feasibility Condition in Political Theory," *The Journal of Political Philosophy*, 1, 27-40.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Belknap Press.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, T.M. 1998. *What We Owe To Each Other*. Cambridge, M.A.: Bellknap Press of Harvard University Press.
- Schroeder, Mark. 2009. "Means-Ends Coherence, Stringency, and Subjective Reasons," *Philosophical Studies*, 143, 223-48.
- Schroeder, Mark. 2011. "Ought, Agents, and Actions," *The Philosophical Review*, 120, 1-41.
- Sinnott-Armstrong, Walter. 1984. "'Ought'" conversationally implies 'can,'" *Philosophical Review*, 93, 249-61.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.

- Smith, Michael. 1995. "Internal Reasons," *Philosophy and Phenomenological Research*, 55, 110-31.
- Southwood, Nicholas. 2010. *Contractualism and the Foundations of Morality*. Oxford: Oxford University Press.
- Southwood, Nicholas. 2015. "The Relevance of Human Nature," *Journal of Ethics and Social Philosophy*, 9, 1-8.
- Southwood, Nicholas. 2016. "'The Thing to Do' Implies 'Can,'" *Noûs*, 50, 61-72.
- Southwood, Nicholas. Forthcoming. "Constructivism About Reasons," *The Oxford Handbook of Reasons and Normativity*, ed. D. Star. Oxford: Oxford University Press.
- Southwood, Nicholas and Lina Eriksson. 2011. "Norms and Conventions," *Philosophical Explorations*, 14, 195-217.
- Southwood, Nicholas and David Wiens. Forthcoming. "'Actual' Does Not Imply 'Feasible,'" *Philosophical Studies*.
- Stocker, Michael. 1971. "'Ought' and 'can,'" *Australasian Journal of Philosophy*, 49, 303-16.
- Swift, Adam. 2008. "The Value of Philosophy in Nonideal Circumstances," *Social Theory and Practice*, 34, 363-87.
- Velleman, David. 2000. *The Possibility of Practical Reason*. Ann Arbor: University of Michigan Press.
- Vihvelin, Kadri. 2004. "Free Will Demystified: A Dispositional Account," *Philosophical Topics*, 32, 427-50
- Wallace, R. Jay. 2001. 'Normativity, Commitment, and Instrumental Reason', *Philosophers' Imprint*, 1/3 (December).
- Wallace, R. Jay. 2014. "Practical Reason," *Stanford Encyclopedia of Philosophy*, available at <http://plato.stanford.edu/entries/practical-reason>.
- Watson, Gary. 1982. "Free Agency," *Free Will*, ed. Gary Watson. Oxford: Oxford University Press.
- Wiens, David. 2013. "The Demands of Justice, Feasible Alternatives, and the Need for Causal Analysis," *Ethical Theory and Moral Practice* 16, 325-338.
- Wiens, David. 2014. "'Going Evaluative' to Save Justice from Feasibility: A Pyrrhic Victory," *Philosophical Quarterly*, 64, 301-07.
- Wiens, David. 2015. "Political Ideals and the Feasibility Frontier," *Economics and Philosophy*.
- Wolf, Susan. 1990. *Freedom Within Reason*. New York: Oxford University Press.
- Zimmerman, Michael J. 2006. "Moral Luck: A Partial Map," *Canadian Journal of Philosophy*, 36, 585-608.