

Forthcoming in *Philosophy Compass*

The feasibility issue¹

Nicholas Southwood (ANU)

We expend considerable time and energy wondering, reasoning, conjecturing, and disagreeing, about what is feasible and infeasible. A Google search for the word “feasible” yields almost 100 million hits. And, of course, this doesn’t account for our private conversations, or what happens inside our heads, or when we deploy the concept of feasibility by using different words: “possible” and “impossible;” “viable” and “non-viable;” “achievable” and “unachievable;” “practicable” and “impracticable;” what we “can” and “can’t” do; and so on.

The prevalence of feasibility is hardly surprising. It is commonly taken for granted that questions of feasibility are highly relevant to our thinking about *normative* questions.² This is perhaps especially true of our normative thinking about *politics*. Thus, we typically assume that in thinking about whether states should introduce a basic income, or open their borders, or start a war, it matters what courses of action, policies, and institutional arrangements are feasible and infeasible. But *what* is it that matters exactly and *how*?

These are difficult and important questions. We know that an inadequate understanding of, and sensitivity to, questions of feasibility may result in normative thinking that is naïve, ineffective, even catastrophic. The twentieth century's failed experiments in Marxism are reminders of what can happen when normative thinking becomes decoupled from credible assumptions about institutional constraints and human motivation. Yet, we also know that normative thinking can be improperly inhibited by misguided ideas about feasibility, leading us to settle for less than we should; to make apologies for the status quo; to let political agents too easily off the

¹ I am very grateful to David Estlund, Matthew Lindauer, Philip Pettit, Laura Valentini, David Wiens, and two anonymous referees for penetrating and detailed written comments on earlier versions of this article and to many other friends and colleagues – too numerous to list here – for invaluable discussion. Research for the article was supported by an Australian Research Council-funded Future Fellowship (FT160100409).

² Robert Goodin and Philip Pettit go further, holding that “questions about what can feasibly be achieved in a certain area *are just as central to normative concerns* as questions about what is desirable in that area” (Goodin and Pettit 1995, p. 1: italics added).

hook; to embrace a cynical realpolitik; to prematurely shut down debate; and to unduly circumscribe ambition and curtail imagination.

The current article aims to offer an opinionated introduction to and characterisation of the *feasibility issue*, as I shall call it, as well as an account of where work on it is currently at. Until very recently, the feasibility issue had received virtually no sustained attention. However, the last few years have witnessed an explosion of interest among political philosophers in the feasibility issue and the emergence of an important and very lively research program.³ This has achieved significant progress. Yet, as we shall see, a number of crucial questions remain unanswered – and indeed largely unasked.

The paper is in three main sections. Section 1 addresses the question of feasibility's nature. Section 2 addresses the question of its role. To anticipate: I shall suggest that there are two quite different kinds of role questions that may be at play, though they are easily (and often) run together: a question about feasibility's *normative significance*; and a question about feasibility's *proper use* in informing our normative thinking. Section 3 considers how the feasibility issue, thus construed, differs from certain other issues in the neighbourhood: the demandingness issue; the issue of whether "ought" implies "can;" and the "ideal versus non-ideal theory" issue.

I. The nature of feasibility

The concept of feasibility pervades our normative talk and thought. But how should it be understood? What is it that we are wondering, reasoning, and disagreeing about when we wonder, reason, and disagree about what is feasible and infeasible? I shall assume that we are looking for an account of feasibility that can make sense of core features of our *actual practices* of making feasibility claims and treating them as relevant to normative questions.⁴

³ For a sample, see Raikka 1998; Cohen 2008; Mason 2004; Brennan and Pettit 2005; Brennan and Southwood 2007; Southwood 2015; 2016a; Southwood and Wiens 2016; Estlund 2007, ch. 14; 2011; 2014; Swift 2008; Gilabert 2011; 2017; Gilabert and Lawford-Smith 2012; Lawford-Smith 2012; 2013a; Hamlin and Stemplowska 2012; Miller 2013; Gheaus 2013; Wiens 2013; 2014; 2015.

⁴ A different approach would be simply to treat "feasibility" as a place-holder for whatever is *in fact* relevant to normative questions, irrespective of our actual practices: say, whatever kind of modal notion is in fact a constraint on truths about what we ought to do. At the same time, our approach does not mean that an account of feasibility must be hostage to *ordinary usage* of the terms "feasible" and "infeasible" (and related terms), departures from which may be justified so long as we can show that such usage is at odds with the way feasibility would have to be in order to play certain key roles that our practices presuppose. I am very grateful to Laura Valentini for forcing me to clarify what sort of account of

A. Cost-based accounts

Three main accounts of feasibility have been proposed. *Cost-based accounts* hold that feasibility is to be understood in terms of what is achievable without undue costs (Raikka 1998; Buchanan 2004; Miller 2013).⁵ Thus, it is not feasible for the United States to eliminate global poverty, or have completely open borders, or transform itself into a communist Utopia insofar as doing so would be unduly (e.g. morally and/or prudentially) costly.

Yet the question of whether a proposal is infeasible and the question of whether its realization would be (even extremely) costly just seem to be different questions.⁶ Moreover, conflating these questions means that cost-based accounts have difficulties explaining certain key ways that we seem to treat feasibility as relevant to our normative thinking. We typically assume that feasibility is special inasmuch as it is a precondition for appropriate consideration in practical deliberation. Where an idea is genuinely infeasible, we assume that it is not appropriate to treat it as among the options for potential consideration in practical deliberation. By contrast, this is not necessarily true of an idea that we take to be unachievable without undue costs – say, in cases where alternatives to the proposal are unachievable without even greater costs or perhaps unachievable simpliciter.⁷

feasibility I am assuming we are looking for and what sort of evidence would bear upon the question of what is the correct account.

⁵ Cost-based accounts come in two versions: objective and subjective versions. *Objective* cost-based accounts hold that feasibility is to be understood in terms of what is achievable without *objective* undue costs (see Raikka 1998, Buchanan 2004, p. 61). Objective costs of an act are consequences of the act that are, as a matter of fact, (e.g. morally and/or prudentially) disvaluable. Objective undue costs are objective costs that are as a matter of fact disproportionate or unreasonable or inappropriate. *Subjective* cost-based accounts hold that feasibility is to be understood in terms of what is achievable without *subjective* undue costs (see Miller 2013). An act is unduly subjectively costly for the agent if it requires the agent to act in ways that are at odds with her fundamental values and convictions.

⁶ A referee made the intriguing suggestion that any evaluative analysis of feasibility such as the cost-based account seems to be vulnerable to a kind of “open question” argument of the kind associated with G.E. Moore. That is, it always seems appropriate to respond: “So, it would be extremely costly/undesirable to X, but is it infeasible to X?”

⁷ Cost-based analyses also face other problems. One concerns what it means for a proposal to be “achievable without undue costs?” It had better not mean that it is “feasible to achieve it without undue costs” otherwise cost-based accounts will be viciously circular. This suggests that it will have to be unpacked in some other way (say, in probabilistic or possibilistic terms) and, hence, that cost-based accounts are not an alternative account of feasibility so much as a particular version of a probability-based or possibility-based account. Another problem is that while they may seem to square well with the way ordinary individuals make feasibility judgements, other linguistic data seems fatal to them. Lindauer and Southwood ms found that individuals’ propensity to resist ascribing feasibility where there are serious costs can be “cancelled” by making the costs explicit. This strongly suggests that the idea of undue

B. Probability-based accounts

Next, there are *probability-based accounts* of feasibility. *Simple* probability-based accounts hold that the feasibility of an agent's realizing a state of affairs is a matter of the probability that the agent will realize that state of affairs. Such accounts can't be right (Brennan and Southwood 2007). It is virtually certain that Australia won't invade Tuvalu, that Vladimir Putin won't refrain from running for a fourth term as Russian President, and that I won't do a chicken-dance in front of my boss (Estlund 2008, p. 13). Yet each of these things is perfectly feasible.

More plausible are *conditional* or *dispositional* probability-based accounts. I shall focus here on conditional accounts, which understand feasibility in terms of certain kinds of subjunctive conditionals.⁸ For example, Geoffrey Brennan and Nicholas Southwood (2007) proposed that it is feasible for an agent or set of agents A to realise a state of affairs x iff A would be reasonably likely to succeed in realising x if A were to try to realise x. (Important modifications of and refinements to this sort of conditional account have been proposed by e.g. Lawford-Smith 2012, Gilabert and Lawford-Smith 2012, Gilabert 2017, Stemplowska 2016.) On this view, it is feasible for the Australian government to process and resettle refugees in Australia rather than offshore inasmuch as and because, even though the Australian government is strongly disposed not to *try* to do this, it is also strongly disposed to succeed insofar as it tries. By contrast, it is infeasible for the Australian government to eradicate the use of recreational drugs since, try as it might, there is simply no reasonable prospect of success.

However, conditional and dispositional accounts are vulnerable to counterexamples involving cases of "volitional incapacity" (Southwood and Gilabert 2016; cf. Lehrer 1968; Wolf 1990) Suppose that Michael Phelps happens to have a pathological phobia of jellyfish in virtue of which he is in some sense "incapable" of trying to save a drowning swimmer at his favourite beach because there are (albeit

costs is playing some non-constitutive (e.g. pragmatic or context-determining) role with regard to our feasibility judgements.

⁸ Dispositional accounts understand feasibility in terms of certain dispositions. For example, Southwood and Gilabert ms suggest that it is feasible for an agent or set of agents A to realise a state of affairs x iff A is disposed to be reasonably likely to succeed in realising x insofar as A tries to x (see also Southwood 2016a). (There are also non-probability-based conditional or dispositional accounts but I shall set these aside in what follows.) The main advantage of dispositional accounts over conditional accounts is that they avoid the problem of so-called "finkish" dispositions – at least on the now virtually universally shared assumption that dispositions are not to be analysed in terms of subjunctive conditionals.

completely harmless) jellyfish in the water. Of course, being a rather good swimmer, he would have no difficulty in saving the swimmer if he *were* to try: that is, if he were somehow to overcome this phobia. But, as a matter of fact, there is simply no prospect of his doing so. Conditional and dispositional analyses would seem to imply that it is feasible for Phelps to save the drowning swimmer. They would also seem to imply that we cannot negate the claim that Phelps ought to save the drowning swimmer by treating feasibility as a constraint on what we ought to do. To many people these implications seem mistaken, though others demur (see Lawford-Smith 2012; Estlund forthcoming).

C. Possibility-based accounts

Finally, *possibility-based* accounts of feasibility understand feasibility in terms of what is *possible* in some relevant sense. The main challenge facing possibility-based accounts is to identify a notion of possibility that will yield a plausible account of feasibility.⁹

Suppose we understand feasibility in terms of some relatively *unrestricted* notion of possibility such as logical, nomological, or metaphysical possibility. Such an account is vulnerable to false positives. It is not feasible for me to single-handedly solve the Israel-Palestine conflict, eradicate poverty, and persuade Donald Trump to adopt Swedish parental leave policies. But each of these things is perfectly logically, nomologically and metaphysically possible.

Proponents of possibility-based accounts therefore typically understand feasibility in terms of some more *restricted* notion of possibility. But restricted how? The most sophisticated and well-developed account of this kind is due to David Wiens. According to Wiens, the feasibility of our realizing a state of affairs requires that our doing so is possible given “our current stock of all-purpose resources” (Wiens 2015, p. 455). All-purpose resources include, not only economic resources but also “the technological, institutional, motivational (and so on) means we have on hand” (p. 453). Possibility is understood as a binary accessibility relation among possible worlds. Thus, Wiens holds

⁹ Another serious difficulty for possibility-based accounts is that feasibility sometimes seems to come in *degrees* (Gilabert 2009; Gilabert and Lawford-Smith 2012; Lawford-Smith 2013a; Gilabert 2017) whereas possibility doesn't. For example, we may judge that Australia's closing offshore processing centres for asylum-seekers this year is *more feasible* than Australia's closing the centres this week. By contrast, our judgements about possibility do not seem to come in degrees; rather they seem to be purely on/off. Closing the centres this year and/or this week is either possible or not possible (though see Kment 2014, chs. 1 and 2).

that “realizing a state is feasible only if there is at least one world at which the state is realized that is circumstantially accessible from the actual world; realizing the state is otherwise infeasible” (p. 458).¹⁰

Appealing as it is, Wiens’ account appears to be incomplete in at least two respects. First, Wiens suggests at one point that certain resources are to be understood or explained in terms of individuals’ abilities or capacities, which are clearly related to, and themselves raise many of the same puzzles as, feasibility itself. “[M]otivational constraints are defined by (the limits of) our motivational resources, namely, human agents’ *capacity* to be motivated to behave in certain ways” (p. 453: italics added).¹¹ Second, Wiens only (officially) postulates a necessary condition for feasibility. If it’s also supposed to be sufficient, then it would seem to be vulnerable to counterexamples involving cases of *counterfactual flukes*: i.e. states of affairs that would involve flukes if they were to be realized (Southwood and Wiens 2016). For example, there is obviously “at least one world at which” I win 10 different lotteries in a single day, “that is circumstantially accessible from the actual world.” This is a problem since Wiens accepts that the feasibility of an agent’s realising a state of affairs requires that the state of affairs not be counterfactually fluke.

D. The problem of groups and institutions

Each of the three main existing accounts of feasibility therefore faces serious difficulties.¹² But even if these difficulties could be overcome, there is a further problem. We appear to assume that feasibility is especially important when it comes to our normative thinking about *politics*.¹³ In politics the central units are, of course, *groups* (or collectives) rather than individuals: legislatures, courts, states, coalitions of states,

¹⁰ Jon Elster (1978) also proposed understanding feasibility as restricted possibility. I am grateful to a referee for drawing my attention to this.

¹¹ A referee has suggested that including motivational constraints in particular may also move Wiens’ account in the direction of conditional accounts in ways that may make it vulnerable to some of the same worries.

¹² Of course, we should remain open to the possibility that some more satisfactory refinement of one of these existing accounts will appear. And it must be emphasized that progress can be (and has been) made in understanding, as it were, the *contours* of the concept of feasibility without giving an account of it. For example, work is ongoing regarding: how if at all feasibility is normative (Lindauer and Southwood ms); how to think about the diachronic character of feasibility (Gilbert 2017); the relation between feasibility and actuality (Southwood and Wiens 2016); and so on. Indeed, investigating the contours of feasibility may even point towards a better account of feasibility.

¹³ For example, Otto Von Bismarck memorably remarked that politics is “the art of the possible, the attainable – the art of the next best.”

political parties, trade unions, intergovernmental and non-governmental organisations, advocacy groups, corporations, social movements, and so on (Lawford-Smith 2012). While some of these groups plausibly constitute group agents (Pettit and List 2012), others don't. Moreover, many important normative political questions concern which sorts of *institutions* we ought to *have*, or how institutions ought to *be*. Yet existing accounts of feasibility are poorly equipped to accommodate either consideration. That's because they treat feasibility as a so-called *agentive modal*.¹⁴ Agentive modals (such as ability and capacity) presuppose an agent and apply to actions. (For example, the ability to speak French presupposes an agent who has or lacks the ability and applies to the action of speaking French.)

Treating feasibility as an agentive modal means, first, that claims about what is feasible for a group that does not constitute an agent have a false presupposition and are, therefore, either all false or neither true nor false – unless they are reinterpreted as claims about what is feasible for (individual) agents (see Lawford-Smith 2012; Collins 2013). Perhaps that is right. But it's worth noting that both options have non-trivial costs. The false presupposition option has drastically revisionary implications for our ordinary thinking about feasibility in politics. Many of the most important feasibility questions in politics turn precisely on whether there is any prospect of turning a disunified collection of individuals into something more unified that meets conditions for group agency, or of reaching an "overlapping consensus" (Rawls 1993) among disunited individual agents that avoids the need for group agency. Yet the reinterpretative strategy arguably neglects key relational impediments to feasibility. Even if it is feasible for every individual member to do her part in a joint activity, it is a further question whether it is feasible for members of the group to *work together* in ways that are required for successful group action.¹⁵

Treating feasibility as an agentive modal also runs into difficulties when it comes to accommodating claims about whether it is feasible for, say, states to have certain

¹⁴ This is perhaps less obviously true of the cost-based account and Wiens' resource account. Regarding the cost-based account: I take it that "achievable without undue costs" is supposed to mean "achievable *by some agent* without undue costs." Regarding Wiens' account: Wiens is quite explicit that he understands feasibility as a property of "realising states of affairs" (see Wiens 2015, p. 459) and that realising a state of affairs is a kind of action.

¹⁵ A referee raised the interesting objection that 1) feasibility is a property of doings, 2) which require doers and 3) that doers are just agents. However, this argument strikes me as pretty questionable. I myself concede (2) but deny (1) and (3). I deny (1) because it seems to me that feasibility may be a property of how we are as opposed to what we do. I deny (3) because it seems to me that agents are *special* kinds of doers.

institutions. Once again, many of the most important questions about feasibility in politics and much of the most significant social science that bears on them appear to involve precisely questions of this kind. Consider, for example, social scientific work on whether certain institutions are internally consistent (think of Arrow's (1950) "impossibility theorem"); or diachronically stable (think of Axelrod's (1984) claim that indefinitely iterated "tit-for-tat" represents the sole evolutionarily successful cooperative strategy over time); or workable in practice (think of Dahl's (1970, pp. 67-8) famous "back-of-the-envelope" refutation of direct democracy); or compatible with other institutions (think of Putnam's (1993) finding that democracy requires special social norms). In each case, this work seems to bear directly on the question of which institutions it is feasible for states to have, or how the institutions within the state ought to be. At least on the face of it, such questions do not seem to involve any category mistake. But nor is it plausible to reinterpret them in a way that involves an action: say, the action of realizing or maintaining the institutional arrangements. For example, the question of whether democracy is feasible in Equatorial Guinea seems to be a different question from, and not settled by, the question of whether it is feasible for any individual or group agent to realise or maintain democracy in Equatorial Guinea. (Perhaps the only realistic prospect of democracy in Equatorial Guinea would involve gradual evolution of core democratic institutions that are then maintenance by a set of social norms.)

In sum, and while this is by no means uncontroversial, if we want an account of feasibility that can accommodate the role of non-agentive groups and institutions, it may be that we need a very different sort of account of feasibility: one that does not treat it as an agentive modal at all.¹⁶

II. The role of feasibility

Let us suppose that, in spite of lacking a satisfactory account of it, we nonetheless have a tolerable implicit grasp of what feasibility involves. What is its appropriate role with regard to our normative thinking? How (if at all) is the idea of feasibility relevant to our thinking about normative questions?

¹⁶ A number of scholars are beginning to work on this task. See, for example, Estlund forthcoming and Southwood ms2.

A. The normative significance of feasibility

As flagged above, I want to suggest that there are, in fact, two quite different questions here: a question about feasibility's normative significance; and a question about feasibility's proper use. Let's begin with the question of feasibility's normative significance. This is the question of whether (and if so how) truths about feasibility have implications for normative questions: what we ought and have reason to do; what justice requires of us; what we are morally permitted, forbidden and required to do; and so on. There are many different ways in which truths about feasibility might potentially be normatively significant. For example, the feasibility of an act might be a *reason* to perform it;¹⁷ infeasibility might be an *excuse* (for not doing what one ought, or what justice requires, or whatever);¹⁸ the feasibility of some *other* act may be relevant to whether I ought to perform an act;¹⁹ the feasibility of some *response* to an act might be relevant to whether we ought to perform the act;²⁰ and feasibility might sometimes be *valuable* such that we have reason (or justice or morality requires us) to *change* or *maintain* what is feasible.²¹

However, the kind of normative significance that has received by far the most attention (and that I shall focus on in what follows) is a matter of whether feasibility is a *constraint* on the truth of some class of normative claims. Many appealing claims about the organisation of social and political life might appear to make infeasible demands. For example, we might think that western states ought to create a world state, or implement completely open borders, or institute an enforceable system of fair trade, or eliminate severe poverty, or drastically curtail global carbon emissions (before it's too late). At least some of these claims might seem to make infeasible demands. If feasibility

¹⁷ This is suggested by Pablo Gilabert and Holly Lawford-Smith (2012), who propose that if it is more feasible for an agent A to X than to Y, then this constitutes a reason (albeit a defeasible reason) for A to X. But now imagine that it is slightly more feasible for a ghastly dictatorial regime to carry out one of two equally horrendous. Do we really want to say that this provides the regime with a reason to carry out the genocide?

¹⁸ For discussion see Estlund 2011, pp. 230-35. However, a possible worry arises in cases where agents are *responsible* for its being infeasible for them to act in certain ways. For example, suppose that I intentionally make it impossible to return your prized first edition of Isaac Newton's *Principia Mathematica* by burning it (Southwood 2016b). (To be sure, we might try to avoid this worry by indexing feasibility judgements more carefully to times. Thanks to a referee for pointing this out.)

¹⁹ For example, it may be relevant to the question of whether I ought to *accept* a refereeing assignment that it is feasible for me to *complete* it in a timely fashion (Jackson and Pargetter 1986; Estlund 2011).

²⁰ For example, what we might call the *feasible enforcement thesis* holds that we ought to implement a rule only if it is feasible to enforce the rule.

²¹ For example, it is plausible to suppose that justice requires governments to make it (more) feasible for members of socially disadvantaged groups to access healthcare and education. For discussion, see Gilabert 2009 and 2017.

is a constraint on the validity of normative claims of this sort, then it follows that they are false just insofar as and because they make infeasible demands.

Is feasibility a constraint (on, say, claims about what we ought to do, or what justice requires us to do, or whatever)? This is a difficult issue on which many of us feel deeply divided (Southwood 2016a). On the one hand, there seems something undeniably *wrong* with normative claims that make infeasible demands; to proffer such claims seems to involve a kind of *wishful-thinking* (Brennan and Southwood 2007; Wiens 2015; Gilabert 2017; cf. Lawford-Smith 2013a). As Philip Pettit nicely puts it, normative theories that encompass such claims “often seem like moral fantasies: manuals for how God ought to have ordained the order of things ... rather than real-world manifestos for what the state should do in regulating the affairs of its citizens” (Pettit 2012, p. 126; see also Brock 2009, p. 4). On the other hand, there also often seems something undeniably *right* about normative claims that make infeasible demands; and to deny this seems to involve treating agents with undue *lenience* (Estlund 2011; Cohen 2007). This seems especially true in cases where we are at least in part *responsible* for the circumstances or character traits in virtue of which it is infeasible for us to do something. In such cases, to insist that feasibility is a constraint on what justice requires us to do often appears to involve letting us too easily off the hook.

How should we make sense of these conflicting reactions? It is natural to suppose that one or other of our reactions involves an error. But which one and why are we disposed to make it? One view holds that feasibility is indeed a constraint and the error consists in thinking that it isn't. The challenge for this view is to say what exactly the error consists in and why we are we disposed to make it. To meet this challenge, it must be shown that we are confusing the correct idea that feasibility is a constraint on claims about what we ought to do (or what justice requires us to do) with some other incorrect idea. There are a number of different candidates: say, that feasibility is a constraint on certain *evaluative* claims (claims about what is good or what is required for perfect justice to obtain); or a constraint on normative claims concerning how we ought to *be*; or normative claims concerning what we ought to *try* to do; or a constraint on how we ought to *deliberate*.

The other view is that feasibility *isn't* a constraint on the relevant class of normative claims and the error consists in thinking that it is. Again, the challenge is to

explain exactly what the error involves and why we make it. This requires showing that we are confusing the incorrect idea that feasibility is a constraint with some other correct idea. Again, there are a number of different candidates: say, that feasibility is a *pragmatic* condition on the relevant normative claims (e.g. that it is conversationally or conventionally implied or presupposed); or a *substantive contributor* to what we ought to do (e.g. a reason for action (see Gilabert and Lawford-Smith 2012); or a constraint on some *other* class of normative claims (e.g. claims about what we ought to do *all-things-considered* (Estlund 2011, p. 226; Cohen 2008, pp. 84-5) or claims about *the thing to do* (Southwood 2016b)).

Alternatively, we might hold that we are *right* to be ambivalent in the face of normative claims that make infeasible demands. Rather than involving any error, our ambivalence is tracking the fact that one and the same normative claim may be supposed to be fit to be used in different ways, and that feasibility is a constraint on the claim's fitness to be used in some ways but not a constraint on the claim's fitness to be used in other ways (Southwood 2016a).²² The most pressing question facing a view of this kind is to say what are the salient uses to which normative claims are supposed to be fit to be put (for which feasibility is and isn't a constraint, respectively). One natural thought is that the key contrast is between prescriptive and merely evaluative uses (Gilabert 2011, p. 56; Gheaus 2013).²³ However, a view of this kind cannot capture the sense in which denying that, say, justice requires Australia to close its offshore processing centres for asylum-seekers this week seems to amount to letting Australia too easily off the hook. This is supposed to be captured by the merely evaluative interpretation of the claim: namely, that if Australia doesn't close its offshore processing centres for asylum-seekers this week, then (perfect or sufficient) justice will not obtain.

²² We might also defend this view by embracing some kind of *pluralism* about feasibility. While I have considerable sympathy for a certain sort of pluralism about feasibility, I won't discuss this option here since I am assuming that our ambivalent reactions may at least sometimes involve the same notion of feasibility.

²³ For example, Anca Gheaus helpfully notes that the locution "justice requires an agent A to X" can be supposed to be used either prescriptively or merely evaluatively. Consider, for example, the claim that justice requires Australia to close its offshore processing centres for asylum-seekers this week. When this is supposed to be used prescriptively it is to be interpreted as something like "justice requires of Australia that Australia close its offshore processing centres for asylum-seekers this week." When used merely evaluatively it is to be interpreted as something like "Australia's closing its offshore processing centres for asylum-seekers this week is a necessary condition for (perfect or sufficient) justice to obtain." Moreover, Gheaus observes that merely evaluative and prescriptive uses of claims about justice "are logically independent and hence that it is possible for the first to be true in a situation in which the second is not." In particular, the second "may be false in precisely those situations when justice recommends the [infeasible]" (Gheaus 2013, p. 457).

But such an interpretation clearly doesn't capture our sense that Australia is subject to a demand of justice to close its offshore processing centres for asylum-seekers this week; that justice requires this *of* Australia; that insofar as Australia fails to do so, it will have done something *criticisable*.

A different version holds that the salient uses are *deliberative* and *hypological* uses (Southwood 2016a). According to this version, feasibility is indeed a constraint on normative claims insofar as they are supposed to be fit to be used *deliberatively*: that is, in practical deliberation. Practical deliberation involves aiming to settle the question of what to do (Southwood 2016b). Insofar as a normative claim demands what is infeasible, then it is not fit to be used to settle the question of what to do, and insisting otherwise would be to be guilty of wishful-thinking. But feasibility is *not* a constraint on normative claims insofar as they are supposed to be fit to be used *hypologically*: that is, to *criticise* agents insofar as they fail to act in relevant ways. Criticism has the aim of holding others and ourselves accountable. A normative claim that demands what is infeasible may be perfectly fit to be used to hold agents accountable, and insisting otherwise would be to treat agents with undue lenience.²⁴

B. The proper use of feasibility

I have been focusing on the question of feasibility's normative significance. Let us now turn to the question of feasibility's *proper use*. This is the question of what way(s) (if any) it is appropriate and inappropriate for us to *deploy* the concept of feasibility in our normative thinking.

The question of feasibility's proper use and the question of feasibility's normative significance are almost invariably run together. But they are distinct in at least three respects. First, the question of feasibility's proper use concerns the role of *judgements* about what is feasible rather than the role of *truths* about what is feasible. This means that the question of feasibility's proper use calls for attending to *epistemic* considerations. Are our feasibility judgements in good order epistemically speaking?

²⁴ A referee helpfully notes that this assumes that there is a kind of *asymmetry* regarding the deliberative and the hypological. That is, while a) it cannot be the case that you deliberatively ought to X if it is infeasible for you to X, b) you can be criticized for not X-ing even if it wasn't feasible for you to X. This is absolutely correct and to my mind perfectly unproblematic. What would be problematic is if we observed an asymmetry *within the domain of the hypological*. That is, it would be deeply problematic if we were to hold that while a*) it cannot be the case that you hypologically ought to X if it is infeasible for you to X, b) you can be criticized for not X-ing even if it wasn't feasible for you to X. But there is no such asymmetry

One problem is that our feasibility judgements seem to be susceptible to various kinds of *performance errors* (Southwood ms1). They often reflect *conceptual confusions and mistakes*. (For example, we mistakenly judge that X is infeasible because we judge that X is unlikely to happen.) They are often *normatively contaminated* (Lindauer and Southwood ms). (For example, we mistakenly judge that X is infeasible because we judge that X is wrong.) And they involve *substantive slippage*. That is, we judge that X is infeasible because we judge that X* is infeasible. (For example, we judge that a carbon neutral economy is infeasible because we judge that a carbon neutral economy without higher petrol prices is infeasible.)²⁵

Second, the question of feasibility's proper use concerns the *procedural* question of how (if at all) it is appropriate to bring these judgements to bear in the activity of answering normative questions rather than the *substantive* question of how feasibility considerations might constrain or help to determine the normative status of our conduct.²⁶ This means that the question of feasibility's proper role calls for attending to certain particular ways in which we might seek out and integrate feasibility considerations into our normative thinking. One important way in which we appear to use our feasibility judgements is as what I shall call *deliberative agenda-setters*. Insofar as we judge that a proposal is infeasible, then we don't consider whether to pursue it at all (and regard it as inappropriate to do so). Yet the susceptibility of our feasibility judgements to the aforementioned performance errors (conceptual confusion, normative contamination and substantive slippage) should make us at least wary about using our feasibility judgements in this way (Southwood ms1). It is plainly inappropriate to treat, say, the judgement that a carbon-neutral economy is unlikely to happen, or the judgement that a carbon-neutral economy is very costly, or the judgement that a carbon neutral economy without higher petrol prices is infeasible as licensing not deliberating about whether to have a carbon neutral economy.

Third, the question of feasibility's proper use is implicitly relativized to a *user* ("appropriate to use by whom?") and a *task* ("appropriate to use in doing what?"). (By contrast, it doesn't even make sense to ask for whom and in doing what feasibility is,

²⁵ Moreover, our susceptibility to performance errors of this kind would seem to be reinforced by psychological mechanisms, political considerations, and the fact that the social sciences, arguably our most epistemically reliable means of forming feasibility judgements, are shot through with normative assumptions and questionable assumptions about motivation.

²⁶ To be sure, it is possible that answers to the procedural questions may somehow track answers to the substantive questions (though I am myself sceptical).

say, a constraint on truths about justice.) The user- and task-relativity of the question of feasibility's proper use calls for attending to the very different kinds of users (individuals and collectives; agents, patients, observers, and advisers; theorists and practitioners); and very different kinds of tasks (deliberation, advice, and criticism (whether individual or collective); policy-making and political theorising). Moreover, it may very well be that we want to say very different things about whether it is appropriate for different sorts of users to use feasibility judgements in some particular way in the service of different sorts of tasks. For example, we may want to say that it is one thing for the editor of a journal to conclude that he ought not to send a paper to Professor Procrastinate on the grounds that it is not feasible for Procrastinate to complete the review in time and quite another thing for Procrastinate himself to conclude that he ought to decline the invitation on the grounds that it is not feasible (even if his conclusion is correct (cf. Lawford-Smith 2013b, sect. 4)).

These cursory remarks offer nothing more than a glimpse of some of the considerations that arise when we take seriously the question of feasibility's proper use. However, hopefully they suffice to show that it is a question that deserves much more attention than it has received hitherto.

III. The distinctiveness of the feasibility issue

I shall now say something about how the feasibility issue, as I have characterized it, differs from certain other issues in the vicinity: the demandingness issue; the issue of whether "ought" implies "can;" and the "ideal versus non-ideal theory" issue.²⁷

A. Demandingness

Consider, first, the relation between the feasibility issue and the demandingness issue. The demandingness issue concerns how demanding a conception of morality or justice can be, where a conception of morality or justice is demanding to the extent that it

²⁷ We might also wonder about the relation between the feasibility issue and the *political realism* issue. Political realism can be thought of as the view that there is a distinctive kind of political normativity that cannot be understood in terms of moral normativity, or at least the kind of moral normativity that is at play in interpersonal morality (see Rossi and Sleat 2014). However, most of the focus of political realists hitherto has been on elucidating and defending the negative part of this thesis (that political normativity isn't moral normativity) and comparatively little attention has gone into explaining and defending the positive part of the thesis: that is, saying exactly what political normativity is. For this reason, addressing the question of the relation between the feasibility issue and the political realism issue at this point in time strikes me as premature.

requires agents to bear significant costs. This issue has, of course, loomed large in the debate between consequentialists and non-consequentialists. A frequent complaint that non-consequentialists make of consequentialism is that it is overly demanding.

If we understood feasibility in terms of what is achievable without undue costs, then this suggests a very close relationship between the question of feasibility's normative significance and the demandingness issue.²⁸ However, we have also seen that there are good reasons *not* to understand feasibility in those terms. This being so, the issue of feasibility's normative significance and the demandingness issue would seem to be orthogonal issues. Thus, we may hold that feasibility is a constraint on morality and justice and that undemandingness isn't a constraint (morality mustn't make infeasible demands but may impose extremely high costs on agents); or vice versa (morality and justice must not impose overly high costs on agents but may make infeasible demands).

B. "Ought" implies "can"

Next, consider the relation between the feasibility issue and the issue of whether "ought" implies "can." We saw in section II that the majority of existing work concerning the normative significance of feasibility has centred on the question of whether feasibility is a *constraint* on the truth of certain normative claims. Plausibly, this question simply *is* the question of whether some particular interpretation of the principle that "ought" implies "can" is valid (Brennan and Southwood 2007; Gilabert 2017; cf. Southwood 2016a).

At the same time, clearly this does not mean that there is nothing distinctive about the feasibility issue. First, while the idea that feasibility is a constraint has received the most airtime, we saw that there are many other ways that feasibility might potentially be normatively significant. (For example, treating feasibility simply as a constraint is blind to the potential normative significance of *degrees* of feasibility above a salient threshold.) Second, we also saw that feasibility's normative significance does not exhaust the question of its normative role since this also includes the question of its proper use. And third, we saw that even if we focus exclusively on the question of feasibility as a constraint, feasibility seems to be a *special* kind of "can" that raises distinctive problems and puzzles. For example, we saw that it may potentially have

²⁸ They are not identical since the idea of undue costliness is generally understood to encompass moral costs as well as prudential costs or costs to the agent.

application to groups that do not themselves constitute agents and perhaps even to how we are as well as what we do.

C. Ideal versus non-ideal theory

Finally, consider the ideal versus non-ideal theory issue (Valentini 2012). The ideal versus non-ideal theory issue is not well defined. But one key animating concern of participants within the debate, I take it, is to figure out the nature and (relative) importance of, and the relation between, two canonical modes of normative theorising: one that is supposed to be based on unapologetically *unrealistic* assumptions (ideal theory); the other that is supposed to be based only on *realistic* assumptions (non-ideal theory) (see Valentini 2012).²⁹

It might seem that the feasibility issue represents nothing more than this animating concern. In particular, the question of whether feasibility has a significant role to play with regard to our normative thinking is nothing over and above, and will be settled by, the question of whether we recognise a legitimate role for ideal as well as non-ideal theory. Feasibility has a significant role to play so long as, and only so long as, we're doing non-ideal theory. Feasibility does not have a significant role to play so long as we're doing ideal theory. We will say that there is a kind of (legitimate and important) mode of normative thinking for which feasibility does not have a significant role to play so long as (and only so long as) we regard ideal theory as legitimate and important. A number of scholars appear to think about the feasibility issue in precisely these terms (see Wiens 2015, p. 447).

But this is a mistake. The question of whether feasibility has a significant role to play with regard to our normative thinking is simply orthogonal to the question of whether we regard ideal theory as legitimate. Whereas the question of whether ideal theory is legitimate is a question about whether certain *inputs* to normative theorising (that is, the assumptions that we are making about how things are) may be unrealistic in the sense of being at odds with the way things actually are, the question of whether normative thinking must be appropriately sensitive to feasibility is a question about

²⁹ Laura Valentini (2012) contends that one thing that is sometimes *meant* by “the ideal versus non-ideal theory issue” is precisely the feasibility issue. Thus construed, there is obviously no interesting question about the relation between the feasibility issue and the ideal versus non-ideal theory issue. The relation is simply identity.

whether the *outputs* of normative thinking (e.g. normative principles) may make infeasible demands.

To see this, notice, first and most obviously, that it is perfectly coherent to hold that a) ideal theorising (understood as normative theorising that involves making all kinds of idealisations or unrealistic assumptions) is valid and important but that b) any kind of normative theorising (including ideal normative theorising) is subject to a feasibility constraint and, hence, invalid insofar as it makes infeasible demands. This was John Rawls' (1971) view. Perhaps less obviously, it is also perfectly coherent to hold that a') ideal theory is entirely illegitimate or irrelevant (only non-ideal theory, or theorising that restricts itself to realistic assumptions, is legitimate) but that b') feasibility needn't play a significant role with regard to such theorising.

How about if we treat assumptions about feasibility as *themselves among the inputs*? Thus construed, ideal theorising would involve thinking about what we ought to do, or how things ought to be, on the false assumption that certain things that are, as a matter of fact, infeasible are perfectly feasible (cf. Cohen 2007, p.p. 250-54). Unfortunately, understanding the feasibility issue in these terms would be to badly mischaracterise what is at issue. On the one side, arch feasibility sceptics – those who are sceptical that feasibility has any normative role to play at all – are going to say that ideal and non-ideal theory will issue in identical verdicts. That's because they hold facts about what is feasible make no difference to what we ought to do and how things ought to be. On the other side, arch feasibility enthusiasts – those who regard *any* normative claims as subject to a strong feasibility constraint – needn't find anything objectionable about ideal theorising, thus construed, so long as the principles that it issues in are understood as having a conditional form; they tell us, not what we ought to do, but what we ought to do if things were otherwise than they are.

IV. Conclusion

It is time to conclude. My aim has been to offer a characterisation of the feasibility issue that involves clearly distinguishing certain questions that have been run together and that can illuminate what is distinctive and important (and indeed distinctively important) about it; and to describe what I take to be the current state of play. While significant progress has been made, plenty of important work remains to be done.

References

- Arrow, Kenneth. 1950. "A Difficulty in the Concept of Social Welfare," *Journal of Political Economy*, 58, 328-46.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. Basic Books.
- Brennan, Geoffrey. 2013. "Feasibility in Optimizing Ethics," *Social Philosophy and Policy* 30, 314-329.
- 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.
- Brock, Gillian. 2009. *Global Justice: A Cosmopolitan Account*. Oxford: Oxford University Press.
- Broome, John. 2013. *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
- Buchanan, Alan. 2004. *Justice, Legitimacy, and Self-Determination: Moral Foundations for International Law*. New York: Oxford University Press.
- Cohen, G. A. 2008. *Rescuing Justice and Equality*. Cambridge, Mass.: Harvard University Press.
- Collins, Stephanie. 2013. "Collectives' Duties and Collectivisation Duties," *Australasian Journal of Philosophy*, 91, 231-48.
- Dahl, Robert. 1970. *After the Revolution? Authority in a Good Society*. Yale University Press.
- Elster, Jon. 1978. *Logic and Society*. Oxford: Wiley.
- 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.
- Estlund, David. Forthcoming. *Utopophobia: Political Philosophy Beyond the Feasible*. Princeton, NJ: Princeton University Press.
- Gheaus, Anca. 2013. "The Feasibility Constraint on the Concept of Justice," *Philosophical Quarterly*, 63, 445-64.

- Gilabert, Pablo. 2009. "The Feasibility of Basic Socioeconomic Human Rights. A Conceptual Exploration," *The Philosophical Quarterly*, 59, 559-581.
- Gilabert, Pablo. 2011. "Feasibility and Socialism," *The Journal of Political Philosophy*, 19, 52-63.
- Gilabert, Pablo. 2017. "Justice and Feasibility: A Dynamic Approach," *Political Utopias: Contemporary Debates*, ed. K. Vallier and M. Weber. Oxford: Oxford University Press.
- Gilabert, 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.
- Jackson, Frank and Robert Pargetter. 1986. "Oughts, Options, and Actualism," *The Philosophical Review*, 95, 233-55.
- Kment. Boris. 2014. *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.
- Lawford-Smith, Holly. 2012. "The Feasibility of Collectives' Action," *The Australasian Journal of Philosophy*, 90, 453-67.
- Lawford-Smith, Holly. 2013a. "Understanding Political Feasibility," *Journal of Political Philosophy*, 21, 243-59.
- Lawford-Smith, Holly. 2013b. "Non-ideal Accessibility," *Ethical Theory and Moral Practice*, 16, 653-69.
- Lehrer, Keith, 1968. "Cans without Ifs," *Analysis*, 29, 29-32.
- Mason, Andrew. 2004. "Just Constraints," *British Journal of Political Science*, 34, 251-68.
- Miller, David. 2013. *Justice for Earthlings: Essays in Political Philosophy*. Cambridge University Press.
- Putnam, Robert. 1993. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton University Press.
- Nagel, Thomas. 1991. *Equality and Partiality*. Oxford and New York: Oxford University Press.
- Pettit, Philip. 2012. *On the People's Terms*. Cambridge: Cambridge University Press.

- Pettit, Philip and Christian List. 2012. *Group Agents: the Possibility, Design and Status of Corporate Agents*. Oxford: Oxford University Press.
- Raikka, Juha. 1998. "The Feasibility Condition in Political Theory," *The Journal of Political Philosophy*, 1, 27-40.
- Rawls, John. 1993. *Political Liberalism*. Cambridge, MA: Belknap Press.
- Rossi, Enzo and Matt Sleat. 2014. "Realism in normative political theory," *Philosophy Compass*, 9, 689-701.
- Schroeder, Mark. 2011. "Ought, Agents, and Actions," *The Philosophical Review*, 120, 1-41.
- Southwood, Nicholas. 2015. "The Relevance of Human Nature," *Journal of Ethics and Social Philosophy*, 9, 1-8.
- Southwood, Nicholas. 2016a. "Does 'ought' imply 'feasible'?" *Philosophy & Public Affairs*, 44, 7-45.
- Southwood, Nicholas. 2016b. "'The Thing to Do' Implies 'Can'", *Noûs*, 50, 61-72.
- Southwood, Nicholas. ms1. "Feasibility in deliberation," unpublished paper.
- Southwood, Nicholas. ms2. "Feasibility for collectives," unpublished paper.
- Southwood, Nicholas and Pablo Gilabert. 2016. "Ability and Volitional Incapacity," *Journal of Ethics and Social Philosophy*, 10, 1-7.
- Southwood, Nicholas and Pablo Gilabert. ms. "A dispositional analysis of feasibility," unpublished paper.
- Southwood, Nicholas and Matthew Lindauer ms. "Feasibility and the problem of normative contamination," unpublished paper.
- Southwood, Nicholas and David Wiens. 2016. "'Actual' Does Not Imply 'Feasible,'" *Philosophical Studies*, 173, 3037-60.
- Stemplowska, Zofia. 2016. "Feasibility: individual and collective," *Social Philosophy and Policy*, 33, 273-91.
- Swift, Adam. 2008. "The Value of Philosophy in Nonideal Circumstances," *Social Theory and Practice*, 34, 363-87.
- Valentini, Laura. 2012. "Ideal versus non-ideal theory: a conceptual map," *Philosophy Compass*, 7, 654-64.
- Wiens, David. 2013. "The Demands of Justice, Feasible Alternatives, and the Need for Causal Analysis," *Ethical Theory and Moral*

- 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*, 31, 447-77.
- Wolf, Susan. 1990. *Freedom Within Reason*. New York: Oxford University Press.