The Concept of Feasibility: A Multivocal Account

Daniel Guillery[1]

Abstract

A common objection to a proposal or theory in political philosophy is that it is not feasible to realise what it calls for. This is commonly taken to be sufficient to reject a proposal or theory: feasibility, on this common view, operates as a straightforward constraint on moral and political theory, whatever is not feasible is simply ruled out. This paper seeks to understand what we mean when we say that some proposal or outcome is or is not feasible. It will argue that no single binary definition can be given. Rather, there is a whole range of possible specifications of the term 'feasible', each of which selects a range of facts of the world to hold fixed. No single one of these possible specifications, though, is obviously privileged as giving the appropriate understanding of 'feasibility' *tout court*. The upshot of my account of feasibility, then, will be that the common view of feasibility as a straightforward constraint cannot be maintained: in order to reject a moral theory, it will not be sufficient simply to say that it is not feasible.

Introduction

Feasibility (or, rather, infeasibility) is a common basis for criticism of proposals in moral and political philosophy and the practice of politics. Practical or theoretical proposals are often rejected on the grounds that what they call for is not feasible. This form of critique has taken on additional prominence with the development of objections to what has been called 'ideal theory'. One important criticism of theories so labelled has been that it is not feasible to meet the recommendations or requirements they generate.

It is common to think of feasibility as a straightforward constraint on moral and political theory. On such a view, a proposed moral principle, normative political

---

✉ Daniel Guillery
   daniel.guillery@gmail.com

1   The University of Warwick, Coventry, UK

theory or proposed course of political action is simply *ruled out* if what it demands or involves is not feasible. Feasibility, conceived of thus, plays a role analogous to that an '"ought" implies "can"' principle takes ability to play. It cannot be the case that an outcome is one we should seek to bring about, or forms part of a correct political theory, if it is *infeasible*. Call this 'the common view'. Conceiving of feasibility in this way requires a single binary notion of feasibility, according to which outcomes are determinately either feasible or infeasible. In this paper, I aim to argue that there is no *single* privileged binary concept of feasibility that can be assumed to play this role. Our ordinary concept of feasibility (and the best interpretation of feasibility premises in moral arguments), I argue, is *multivocal*; that is, there are many different ways in which it can be made precise (which I call 'specifications'), each with different truth conditions. Not *all* of these can plausibly be thought to be universal constraints on morality, and there is no immediate reason to suppose that any particular one of them is *the* constraint for morality in general.[1] If there is no single, universal feasibility constraint, then it will not do to reject a proposal simply by denying its feasibility: we will need to know more about what that means and what moral significance it has.

Holly Lawford-Smith and Pablo Gilabert have prominently argued that there is a *scalar* notion of feasibility, which does not rule out proposals but contributes to ranking them (Gilabert 2009, 2017; Gilabert and Lawford Smith 2012; Lawford-Smith 2013a, b). If feasibility is scalar, it cannot play the straightforward constraining role that it is taken to play by the common view described above. That part of my conclusion, then, can be arrived at by another route. (For Gilabert and Lawford-Smith, though, 'feasibility' has two senses, the other of which is binary and does play the simple constraining role.) Nevertheless, a scalar account of the concept is consistent with the main feature of the common view that I want to reject. A scalar account can be *univocal*; it can hold that for any outcome, even if there is no answer to the binary question whether it is feasible or not, there is a single determinate answer to the question *how* feasible it is. If 'feasibility' is a univocal scalar concept, it will not act as a simple *constraint*, but it will provide a simple, linear dimension for *ranking* proposals. My claim will not be that feasibility should be understood in scalar terms, but rather that the concept admits of many different specifications, and so, for any given outcome, there is no single determinate answer to the questions whether it is feasible *or* how feasible it is. (I will focus, in what follows, on the many

---

[1] A number of writers in the growing literature on feasibility have attempted to give rigorous accounts of a univocal concept of feasibility that plays something like this straightforward role (see, for instance, Brennan and Southwood 2007; Stemplowska 2016; Wiens 2015). Erman and Möller (2020) have recently argued that feasibility constraints on normative political theories are contextually variant according to the function or purpose of the theory. My argument here will be an attempt to push towards roughly the same endpoint, but from a different direction. Erman and Möller focus quite generally on arguing that the kinds of assumptions and proposals appropriate to normative political theory can vary significantly with the function or context of that theory. My aim here will be to argue that the concept of feasibility itself is not on its own the kind of thing that can play a simple universal constraining role. Another earlier contribution to the literature that may anticipate my argument is Jensen (2009), which could be read as gesturing at a multivocal account.

binary specifications of the concept available, but I do not intend to deny that it may *also* have similarly important scalar senses.)

The primary dialectical import of the paper, then, will be to suggest that things may be more complicated than the common view of feasibility as a simple constraint (or ranking) takes them to be. But there is also a secondary aim, and that is to give a metatheoretical account of what unifies the different binary specifications of 'feasibility'. The different ways of making the term precise are differentiated, I will argue, by how they fill in a single variable: the range of facts held fixed. There is thus a core of what it is to be feasible *given a range of facts held fixed*. The secondary aim of the paper, then, will be to flesh out an account of this core. The unifying account I give is a *possibility*-based account. To claim that some outcome is *feasible* is to make a claim about something being *possible*, but it is not simply to claim that that outcome is possible. Though related, possibility and feasibility are distinct concepts. Feasibility, unlike possibility, has to do with agency.[2] It is about what it is possible for agents to do through intentional action, in a way that my account will make clear, not just what might possibly come about. (It is because of this that feasibility is of such interest to moral and political theory: it is about what is open to us, through our action.)

I will begin by motivating and defending the multivocal account. I will motivate intuitively the idea that the concept of feasibility is susceptible to being specified in multiple ways. Since the primary accounts of feasibility given in the literature are univocal, I will then briefly argue that these fail either to capture our ordinary concept of feasibility or to identify a simple constraint on moral requirement. In the final part of the paper, I will flesh out the details of my account. My aim here will be to show how the different binary specifications of 'feasibility' are unified as specifications of a single, multivocal concept. This will show how feasibility is evaluated, and how it *can* play a constraining role once it is made precise (i.e. once certain variables are filled in), and will demonstrate how a multivocal account can be made complete.

The Multivocal Account

One might assert 'It is not feasible to institute a system of participatory democracy'. Supposing that we know exactly what it would mean for a system of participatory democracy to be instituted (i.e. what such a system is), there are still many different things this statement could mean; it is not clear, without any context, what exactly its truth conditions are. Or that, at any rate, is what I will claim. Put roughly, one thing we might mean is that given people's motivations and preferences being as

---

[2] Holly Lawford-Smith (2013a) makes a similar point (p. 247). That feasibility has to do in some way with agency is a basic fact about the concept I am interested in. This is the ordinary way in which the term is used (the *Oxford English Dictionary* defines it, for instance, as 'capability of being *done*; *practicability*', my emphasis), but if it is occasionally used differently (for instance, synonymously with mere possibility) then it is just being used to mean something different to the concept I am concerned with here.

they currently are, with the parliamentary and electoral systems and the balance of power being as they are, it is not possible to institute participatory democracy. If this is what we mean, the statement is quite plausibly true. However, this is not the only thing that might naturally be understood by the claim; in a different context, we might hold fixed a different range of facts. At the far extreme, we could mean that such a system is not physically, or logically, possible. If we meant *this*, probably our statement would be false. No doubt that is not an interpretation that will be salient in many ordinary contexts. However, there are various other intermediate interpretations, which represent perfectly natural uses of the term 'feasibility'. We could mean, for instance, that, even if people's basic motivations, the power balance and so on are allowed to change, a system of participatory democracy is made impossible by some reasonably deep facts about human nature (of a kind that are not physical laws). There is no immediately obvious reason to suppose that any one of these various possible readings of the claim is privileged over the others as representing the 'proper' use of the term 'feasibility'. It would be a perfectly sensible response to the question 'is it feasible to institute a system of participatory democracy?' to ask 'holding fixed what?'. (Of course, in many ordinary conversational contexts, some particular range of facts to hold fixed will be made salient, implicitly if not explicitly, but if the question is asked with no relevant context, it is not obvious how to answer.)

The data on our ordinary uses of the term 'feasibility', then, suggest that there are at least a number of different ways in which it is used, a number of different ways in which it can be made precise. We are interested in feasibility for purposes of moral philosophy, though. Perhaps this diversity is not present in contexts of moral reasoning. On the common view I mentioned above, there must be some single binary concept of feasibility generally salient in moral contexts. Does our ordinary moral thinking bear this out? I think not: different specifications of 'feasibility' appear relevant when we ask different moral questions.

One obvious kind of moral question we might ask is whether we are, or could be, morally required to bring about some outcome, a system of participatory democracy, say. We naturally want to say that if it is not *feasible* to bring it about, we cannot be morally required to do so. But what does that mean? What is the sense in which it must be feasible in order to be a possible object of moral requirement? One thing to notice is that the proposed object of requirement itself is underspecified. Are we asking whether we are required to bring such a system about *right now*, or in the next year, or merely to act in such a way that will cause it to come about at some point in the future? The feasibility constraints that are relevant appear different when we ask different versions of this question. Facts that constrain what we can be required to do *now* are not all constraints over the long term. But suppose we settle on a timescale. We might ask, for instance, whether we can be required to bring about participatory democracy in our lifetime. It is unclear exactly which feasibility constraints are relevant to this question. One source of uncertainty is empirical uncertainty about how hard or intractable particular feasibility constraints are. But setting that aside, there are still different moral questions we might ask even about this outcome, 'bringing about participatory democracy in our lifetime', for which the relevant feasibility constraints appear to vary.

We might ask, for instance, whether we can be required to *expend our efforts* on bringing about such a system in our lifetime. It is plausibly a corollary of the common view that you cannot be required to expend your efforts on bringing about an outcome O if it is not feasible to bring about O (at least, you cannot be required to expend your efforts in that way *for the sake of* bringing about O).[3] There are many actions we might perform with the aim of bringing about an outcome; for instance, to bring about a system of participatory democracy, we might launch a campaign to build public awareness, we might initiate a programme of legislation in parliament and so on. Obviously, the feasibility of doing these things themselves is relevant to the question whether we can be required to expend our efforts in these ways. But the feasibility of the *outcome* also has some bearing on that question. The common view is correct in that when we ask whether we ought to expend our efforts on bringing about an outcome, the feasibility of that outcome is certainly something we should consider. But when we ask different moral questions about how we should expend our efforts, the relevant specifications of 'feasibility' seem to vary.

Suppose we are a minor political party, with some representation in parliament but only a very small proportion of the seats. We believe that a system of participatory democracy is desirable. One question we may ask ourselves is whether we should expend our efforts on attempting to cobble together a coalition of parliamentary representatives prepared to vote for a project of law that will put us on a path to participatory democracy (and assume as well that this process of coalition building will not itself, unless successful, achieve some part of what is desirable about participatory democracy). Could we be morally required to do so? The common view will tell us that we cannot be if it is not feasible to bring about the outcome of a participatory democratic system. But in what sense must it be feasible in order for us to be required to expend our efforts in this way? There is room for disagreement here, but it is quite plausible that, if holding fixed the current preferences and motivations of the public, the existing parliamentary system and electoral system, as well as the configuration of political parties (as well as basic facts about things such as physics and biology) rules out such a system, that is enough to show that we cannot be required to expend our efforts on finding a parliamentary route to bring it about. For *this* kind of moral question, a relatively restrictive specification of 'feasibility' may constrain us.

But when we ask a different kind of question, this specification of 'feasibility' is less obviously a constraint. Suppose instead we are considering whether we should expend our efforts on attempting to bring about participatory democracy by doing things like advocating for it, attempting to reshape public opinion, supporting minor changes that put us on the right path and so on. There is still

---

[3] Strictly speaking, it could be that even if O is not feasible (and so you cannot be required *to bring it about*) you might nevertheless be required to *try* to bring it about (maybe we could be required to try to eradicate global poverty even if that were not feasible). But this will only be so if *getting closer* to or *approximating* O is both desirable and feasible. If getting closer to, or approximating, O is also not feasible (or is not desirable), then presumably on the common view you cannot be required to expend your efforts on bringing it about (i.e. you cannot be required to expend your efforts on bringing about an infeasible O for the sake of bringing about O).

a valid, ordinary sense in which the outcome (the institution of a system of participatory democracy) is not feasible for us here and now. But it is not so clear that feasibility understood in that way constrains us when asking this new kind of question. It *might* be that we should expend our efforts on attempting to bring about participatory democracy in these ways, and our being morally required to do so is not obviously *ruled out* by the fact that the goal is infeasible when we hold fixed public opinion, the party system and so on. If, on the other hand, it were inconsistent with, say, deep facts about human nature, then our being morally required to expend our efforts in these ways *might* be ruled out. Thus, a different specification of feasibility seems relevant here.

My suggestion is not that these particular specifications of 'feasibility' in fact are the ones relevant to these kinds of practical question, but just that there are multiple different ways of understanding feasibility, and it is not obvious that any one of these constitutes the salient constraint for all practical decisions. (It is plausible that the relevant feasibility constraints for questions about the nature of moral ideals, for instance, are different again, and questions about our *pro tanto* obligations may differ from questions about what we ought to do all things considered.)

If this is right, then, the concept of feasibility is open to specification in multiple ways. The data canvassed about ordinary use and moral reasoning also suggest that what differentiates these specifications is the range of facts of the world we hold fixed. When we hold fixed different sets of facts, we get different answers to feasibility questions. I thus want to propose a framework for thinking about the multiple possible specifications of 'feasibility', according to which each such specification corresponds to what I will call a 'feasibility constraint' (FC). An FC is a selection of which current facts of the world to hold fixed. For each possible FC, there is a possible specification of 'feasibility'. Feasibility, then, is assessed (and feasibility claims have truth conditions) relative to a choice of FC. Often it is obvious from conversational context what specification of 'feasibility' is assumed in talking about feasibility, i.e. a choice of which facts to hold fixed is tacitly assumed (or, less often, one is made explicit). However, this need not always be the case. Sometimes, when we make a feasibility claim we fail to say anything determinate because no choice of FC is understood or specified.

In principle, *any* set of facts could be held fixed and could constitute an FC relative to which a feasibility claim is assessed. But it is clear that certain possible FCs do not give us specifications of 'feasibility' that correspond either to any ordinary uses of the term or to constraints relevant to any practical decisions: for instance, an FC that holds fixed nothing but the laws of physics, or one that holds fixed my current spatial location. I do not aim to offer a principled way of separating out which are the FCs that do correspond to ordinary use in this way and which do not. My aims in this paper are simply to, first, argue that the concept admits of *many different* specifications, and, second, to identify the core of the concept that is common across these different specifications. What I think the foregoing suggests is that there is a fair range of different specifications (with quite different results) that *do* correspond to ordinary use and can be relevant to practical decisions.

Alternative Accounts

Before I go on to flesh out my account in detail, though, I must acknowledge that this appearance of multivocality might be just that, an appearance. It could be that we can give an account of a univocal concept of feasibility capable of explaining away this appearance and showing that these seemingly disparate specifications of 'feasibility' are in fact mere applications of a single univocal concept to different contexts. The leading accounts in the existing literature are univocal (or 'bivocal', in the case of Gilabert and Lawford-Smith), and presumably these accounts are intended to do just that. Thus, in this section I will briefly argue that the two main univocal contenders do not offer a good account either of our ordinary concept of feasibility or its role in moral reasoning. Our intuitions pull us in more than one direction, and these univocal accounts, in forcing us to go one way across the board, have implausible implications.

Conditional Probability Account

What has become the dominant account of feasibility in the literature analyses the concept in terms of conditional probability.[4] This was first proposed by Brennan and Southwood, who rejected both a simple possibility account as well as a simple probability account. First, they point out that there are many things that are logically or nomologically possible that do not qualify as feasible, such as a medical ignoramus performing a neurological operation for which they lack the relevant expertise. These things, they think, are not feasible because, though possible, they are not probable. It is *possible* that by sheer luck the medical ignoramus could perform the exact correct sequence of movements to perform a neurological operation. However, this is extremely unlikely, and we do not want to say that it is *feasible*. On the other hand, feasibility cannot simply amount to probability, because although it may be very improbable that a parent will get out of bed at the weekend to watch their daughter's hockey games, it is clearly not the case that this is thereby infeasible. The conditional probability of their going *if they tried*, however, is presumably much better. Thus, they say, feasibility should be understood 'in terms of reasonable probability of success conditional on trying' (Brennan and Southwood 2007, pp. 9–10). If we flesh this out according to a standard (Lewisian) account of conditionals, to say that O is feasible for you means roughly 'there is a closer possible world at which you try to bring about O and have a reasonable probability of success than any at which you try to bring about O and have an insufficient probability of success' (Lewis 1973, pp. 424–425).

---

[4] This is the account at the core of Gilabert and Lawford-Smith (2012), though they treat conditional probability as scalar. The cases I will discuss below in which the conditional probability account wrongly classes outcomes as feasible or infeasible are also cases in which the *scalar* version of the account wrongly classes outcomes as quite or very feasible or infeasible. Nicholas Southwood (2016) defends a *dispositional* account, which is subtly different, but similar. Most of what I will say here should apply to such an account as well.

A plausible version of the conditional probability account must demand probability conditional on something like *wholehearted* trying, since the closest possible world in which I try to, say, run a mile is not one in which I try very seriously. This gives us the following account:

> (CP) It is feasible for A to bring about O if, and only if, if A *wholeheartedly* tried to bring about O, she would probably bring about O (see Stemplowska 2016).

What is meant by 'wholehearted trying'? It cannot be 'performing the objectively best bundle of actions for O', for then we would end up having to say that it is feasible for the medical ignoramus to perform brain surgery. Thus, it seems we will need to cash out the above as the following:

> (CPb) It is feasible for A to bring about O if, and only if, if A were to pursue whatever is believed to be an effective (or likely to be effective) means to O, she would probably bring about O.

However, there are intuitions that this account does not capture. I focus on what appears to be the most serious problem.

We intuitively want to include at least some motivations as constraints on feasibility but not to always include all motivations as constraints. In Brennan and Southwood's lazy parent case, for instance, there is at least a valid sense in which it *is* feasible for the parent to attend their daughter's hockey games despite their lack of motivation. In other cases, we want to allow certain motivations (such as pathological motivational failures) to count as a constraint. The conditional probability account cannot capture both of these intuitions where a multivocal account can.

Account (CPb) can be read in two ways: 'pursue' in 'pursue whatever means to O are believed to be effective' can be read as a success verb or not. If read as a success verb, there are again two possibilities. First, it could mean '*perform* (successfully) whatever actions are believed to be such that *if they were successfully performed* would be effective for O'. If read in this way, the account seems to exclude all motivations, even extreme pathological ones. If I believe that walking across the plank positioned over the 500 m chasm would be an effective means to cross it then there is an action that I believe to be effective that, if successfully performed, would have a high probability of resulting in my getting across the chasm. But if I suffer from a pathological fear of heights such that I could not bring myself to walk on the plank, I think we would be loath to say that it is feasible for me to cross the chasm.

The second possibility, if we read 'pursue' as a success verb, would take the clause to mean 'successfully perform those actions believed to be effective for O *if attempted*'. This allows an agent's beliefs about their motivations to constrain feasibility to far too great an extent. If the agent believes (correctly or incorrectly) that they will not be motivated to carry through an attempt to bring about O, then O becomes infeasible. If I believe that if I attempt to write a book review, I will soon get distracted and give up, then there is no action I believe to be effective for

O *if attempted*, and so on this interpretation of (CPb) it is not feasible for me to write the review. But there is at least a valid sense in which it is feasible for me, and I may be morally required to do so.

If we do not read 'pursuing' as a success term, we get '*setting out* to perform whatever actions are believed to be effective'. In this case, motivations are ruled in as constraints more or less wholesale, since in the closest possible world in which I try wholeheartedly in *this* sense, it may be that I would not in fact succeed in performing those actions, just because I would not be motivated to carry through. This, too, is implausible: we do not want to say that outcomes are infeasible for me, in the only significant sense, whenever I lack the motivation to carry through sequences of actions that would bring them about. And we certainly do not want to say that lack of motivation always defeats moral requirement (cf. Estlund 2011).

A natural response is that 'pursuing' means neither 'performing' nor 'setting out to perform' but rather 'trying wholeheartedly to perform' but then the question is just postponed ad infinitum. On a multivocal account these seemingly conflicting intuitions are accommodated, since it simply allows that on some specifications of 'feasibility', motivations are constraints and on others they are not. For any given moral inquiry, it is a difficult question which constraints are relevant. I will not attempt to answer that question here. But we have little reason to assume that there must be one single answer valid across all kinds or purposes of inquiry.

Zofia Stemplowska (2016) defends a modified version of the conditional probability account that she thinks deals with the problem of motivational failure. Her suggestion is that when there is a conceivable incentive 'that would bring the agent's motivational state in line with what is needed to perform the action in question', the action is feasible for the agent, whatever may be true about their *actual* motivations, but if there is no such conceivable incentive then the action is infeasible for them (Stemplowska 2016, p. 280).

However, there seems to be at least a sense in which it is *feasible* for me to, say, kill someone I love, even if there is *no* conceivable incentive that would induce me to do so. There may be few such actions, but we can certainly imagine there being some things that we are *so* motivationally committed to not doing that we never will (so long as our motivations remain constant). Nevertheless, it is natural to say that, in some cases at least, we are committed to not doing these actions *despite* their feasibility for us. Stemplowska notes this problem in the case of actions that we are committed to not performing for *moral* reasons. Thus, she revises her definition to:

> Action φ is (more) feasible if there is an incentive I—or had the agent X not seen φ as wrong there would be I—such that, given I, X will try to φ and, given I, X is (more) likely to φ. (Stemplowska 2016, p. 281)

However, this only resolves part of the problem. An agent's robust motivational commitment to not doing an action need not arise for moral reasons. You might be perfectly committed to not doing φ for non-moral reasons, but yet there be no (or little) motivational difficulty in doing φ *if you wanted*. I might resolve to pursue some project (with no particular *moral* value) come what may and be so stubborn or determined that *no* incentive will induce me to do otherwise, even though I *can* easily motivate myself to do otherwise *if* I want to. In addition, it seems perfectly

possible that one could be morally required to do something that one is committed in this way to not doing.

Univocal Restricted Possibility Account

The other main candidate univocal account is that of David Wiens (2015). This, like the account I will develop below, takes feasibility to be a matter of restricted possibility, but, unlike mine, identifies a single relevant constraint. According to Wiens, feasibility should be understood as a possibility consistent with a 'resource stock' (2015). The resource stock defines an accessibility relation on the set of possible worlds and feasibility is a matter of possibility within this accessibility relation: in other words, there being a possible world consistent with the resource stock in which the outcome comes about. Wiens defends this account only as a *necessary* condition for feasibility, since, just like the simple possibility account that Brennan and Southwood rejected, on its own it will allow in many things that should not count as feasible. There is almost certainly a possible world consistent with the 'resource stock' in which our medical ignoramus successfully performs a neurological operation, for example.

However, even as only a necessary condition, Wiens's account is problematic. Wiens defines his accessibility relation (the set of worlds that constitutes the feasible set) in two steps: it includes not only worlds that can be realised given the *actual* resource stock, but also worlds that are realisable given resource stocks that are *attainable* by transformation of the actual resource stock. *Either* this only allows *one* transformation of the resource stock (supposing there is some way of delimiting what counts as a *single* transformation) or it allows multiple iterations.

If we only allow worlds realisable after *one* iteration of transformation of the actual resource stock (whatever that means), then we seem to arbitrarily restrict the feasible set, and we rule out many things that should count as feasible. For instance, suppose that to institute some policy we will need increased economic resources and in order to get these we will need to change public opinion about government spending (but this is quite easily done). On any plausible way of individuating transformations, this presumably will involve multiple successive transformations of the resource stock, but it does seem that there is a sense in which it is feasible. We also may sometimes be morally required to do things involving multiple successive transformations of the resource stock.

On the other hand, if we allow multiple iterations of transformation of the resource stock, the account becomes *very* permissive. There are *very many* (quite unrealistic) outcomes that come about in *some* possible world that is accessible by *some* possible series of transformations of our current resource stock. There is almost certainly, for instance, a *possible* series of transformations of the current resource stock that is consistent with the realisation of a proportional electoral system in the UK. But it is not obviously false to assert that the realisation of such a system is *not* feasible, in at least some valid (and morally relevant) sense.

Existing accounts, then, run into problems in attempting to give a single set of necessary and sufficient conditions for feasibility. They fail to account for all of

the ways in which the concept is used in ordinary language or moral reasoning. It appears that what they miss is that for many outcomes there is both a sense in which they are feasible and a sense in which they are not (and it is not always obvious for a given moral question which of these is the relevant constraint).

Fleshing out the Multivocal Account

Nothing I will present constitutes any sort of conclusive argument for the multivocality of the morally relevant concept of feasibility. So far I have shown that on the face of it, our use of the concept in ordinary language and moral reasoning appears to be multivocal, and that existing univocal accounts encounter problems. That does not rule out that these problems could be overcome, or that another univocal account could be given. But together, they constitute a good preliminary case for returning to the presumption that the initial appearance is correct. To conclude my case for a multivocal account, I want to show how such an account can be fleshed out, and how we can understand what unifies the various different ways in which the concept can be specified.

In this section, then, I will propose a picture of the core of the concept, that is, a picture of what it is to be feasible *given a choice of facts to be held fixed*. Although there is no univocal binary concept of feasibility *tout court*, I think we can give a binary definition of 'feasibility' given a choice of feasibility constraint (FC).[5] I start from the intuitive idea that feasibility is a special form of possibility and then motivate certain modifications in order to deal with problem cases. (Problem cases have pushed some theorists away from possibility-based accounts towards, for instance, the conditional probability account, but once we add some sophistications to a simple possibility account (and go multivocal) we can avoid these sorts of problems.)

If it is asserted that some outcome is feasible for us, it is natural to think about the truth of this as having to do with whether it is *possible* for us to get there, or to bring it about, given certain facts of the world. What is feasible for us, goes the thought, is not simply anything that might possibly come about, but what might possibly be brought about by us.

However, as we have seen, it cannot be as simple as this. Brennan and Southwood's medical ignoramus case showed that simple possibility is not the same as feasibility. Going multivocal, we can also note, does not solve this problem. So long as we grant the medical ignoramus access to surgical tools and ordinary human capacities of bodily movement, it will be *possible* that they could perform the exact sequence of movements involved in successful brain surgery. This will be true whether or not we hold fixed their knowledge of surgery. But if we do hold it fixed, the only possible scenario in which they succeed is one in which they do so by sheer luck. This case motivated Brennan and Southwood to turn to the conditional probability account. But a more sophisticated possibility account need not be subject to

---

[5] As mentioned previously, though, this is not to deny that there could also be a scalar sense given a choice of facts to hold fixed.

such counter-examples, and since it is a more intuitively natural starting place, that is where I will begin.

On the possibility-based account I will present, feasibility is not simply equivalent to possibility. But feasibility can be cashed out *in terms* of possibility. For something to be feasible, as already noted, it needs to be possible for it to come about in a particular way, one that involves agency. For something to be feasible given a set of facts being held fixed is for it to be possible for us to bring it about compatibly with these facts.

Thus, I propose the following definition for binary feasibility given a choice of FC, that is, of which facts to hold fixed, which I will explain and motivate below:

> An outcome O is feasible for an agent X in a context Z given FC *f* if, and only if, it is possible, compatibly with constraint *f*, for X to perform an intentional action that will bring about O (though the action need not be intended to bring about, or contribute to bringing about, O) such that X brings about O *safely* and *competently* (notions to be explained below).[6]

Put more simply, an outcome is feasible for an agent if it is possible (compatibly with the facts held fixed) for them to bring it about through intentional action, with the added proviso to be elucidated below that they must do so safely and competently.

In order to see what is involved in something being possible given some constraint, it may help to think of the constraint as playing a similar role to an accessibility relation in modal logic.[7] An event is possible given an FC if it occurs in some possible world out of a restricted range selected by the choice of constraint (*f*) and context (Z). (The choice of context (Z) is a choice of time and possible world; the choice of FC is a choice of facts, from those that hold in Z, to be held fixed.) The world from which the accessible worlds must be accessible (call this the home world) is selected by Z (it is likely to be the actual world, but need not be: we might want to ask questions about what would be feasible in a counterfactual scenario). The accessible worlds are then restricted to those identical to the home world up until the time of Z, i.e. until the time for which the feasibility question is being asked. This is because, once we have asked about feasibility for an agent in a particular time and possible world, we are interested in the possible *futures* available from that time and world, not counterfactual possibilities in which the past was different. Finally, the FC then restricts the accessible worlds to those in which, after that time, all the facts selected by the chosen feasibility constraint stay fixed. If an outcome is brought about in the right way (directly or indirectly) by X in some possible world out of this restricted range, then it is feasible for X in Z given this FC, i.e. given this specification of 'feasibility'. What this means, in less abstract terms, is that when we choose a range of facts to hold fixed, say the deepest facts of human

---

[6] I will focus here on feasibility for an individual agent. Feasibility for groups raises interesting questions (see, e.g. Stemplowska 2016, pp. 276–277; Lawford Smith 2012). I think this account can be extended to cover group feasibility and feasibility *tout court*, but for lack of space, I will not attempt to do so here. The *definiendum* here is a modified version of that proposed by Gilabert and Lawford-Smith (2012).

[7] See also Wiens (2015, p. 457).

nature along with the laws of physics, biology and so on, an outcome is accessible for me if and only if there is some possible world with our actual history in which those laws and facts of human nature stay constant, in which I bring about the outcome in question.

Although, as I have said, feasibility has to do with what *agents* can bring about through intentional action, my definition does not require that for O to be accessible to X it must be possible for X to bring about O *intentionally*. It allows for the possibility that there may be certain outcomes that it is only possible to achieve when not aiming at them. It could be, for instance, that it is only possible to achieve a state of meditative calm when you are not aiming directly to do so. Similarly, for it to be feasible for you to get a grade of 67 for an exam you take, it does not need to be possible for you to deliberately set out to get a 67 and succeed. The grading may be too fine-grained for that; there may be no way for you to calibrate your performance precisely to get a 67 and not a 68 or 66. Nevertheless, if it is possible for you to intentionally try hard and perform your best in the exam, and your performing well could get you a 67, then it is feasible for you to get a 67. Thus, feasibility does not require that it be possible for you to bring about the outcome in question *intentionally*, but it does require it to be possible for you to do so *through intentional action*.

For O to be feasible for X, though, it is not enough that it be possible for X's intentional action to bring about O. As we have already seen, it is *possible* that a medical ignoramus who sets out to perform brain surgery by trying a random series of movements will successfully perform brain surgery. But, for the outcome to be *feasible*, it needs to be possible for the agent to bring it about *not just by freak luck*. For this reason, I add the requirements of *safety* and *competence*. I borrow these notions from the literature in epistemology, where it is often thought that an account of knowledge must accommodate the intuition that true belief achieved by luck does not count as knowledge (see, for instance, Ichikawa and Steup 2017; Pritchard 2012; Sosa 2007).

Firstly, O's being feasible for X requires that there be a possible action of X's that brings about O *competently*, by which I mean that O is creditable to some sufficient extent to X's relevant *competence*. Ernest Sosa describes a competence as 'a disposition, one with a basis resident in the competent agent, one that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it' (2007, p. 29). I will not attempt to give a full account of what a competence is, but will assume that there is an intuitive notion identified by Sosa's description. A competence is like a skill: some actions that an agent performs manifest competence or skill, while others do not. An experienced archer hitting a target in ordinary conditions and with no intervening factors seems to be an example of the former, while a game-player rolling a six seems to be an example of the latter. It is clear that a medical ignoramus is *not* competent to perform brain surgery, though they *may* be competent to perform the precise sequence of movements that would be needed in a particular instance to perform brain surgery. My account requires that there be a possible world in which X's bringing about O is *sufficiently creditable* to X's competence. An outcome is not ruled out as feasible for me if it is only possible to bring it about with some cooperation from circumstances beyond our control, but it needs to be possible for it to be brought about in a way

that is attributable to some sufficient degree to my competence: if the only possible scenarios in which I bring about O are scenarios in which O is *primarily* attributable to, say, a storm, and I play merely a supporting role, then it is not plausibly feasible for me.[8] (It could be, though, that an outcome like hitting a tricky shot in a game of snooker is something I will fail to do most of the time, but will achieve, say, one time in eight. So long as that possible success is sufficiently attributable to my competence, a competence that enables me to succeed roughly every eight attempts, it can count as feasible for me to hit the shot given about eight attempts.) Note also that my account does not require that X *actually* be competent to bring about O, but rather that it be *possible* for X to be so competent.[9]

I also add a safety requirement. The requirement is that it must be possible for X to bring about O safely.[10] Sosa characterises safety thus: 'A performance is safe if and only if not easily would it then have failed, not easily would it have fallen short of its aim' (2007, p. 25). For a performance to be safe, it needs to be the case that it succeeds not only in the actual world, but also in other nearby worlds, similar to the actual world in certain relevant respects. In the case of feasibility, what is needed is that there be a *possible* world *w* in which X brings about O safely. This means that in all the sufficiently close possible worlds to *w* (the possible world in which X brings about O), in which circumstances are relevantly similar, X succeeds in bringing about O. The addition of a safety requirement on top of the competence requirement is needed because there could be cases where some piece of freak luck makes possible the exercise of a competence. For example, suppose there is a brick wall separating Ella, a competent darts player, from a dartboard. There are some large rocks in the vicinity that, if they were positioned in a particular spot, would allow Ella to climb on top and throw a dart over the wall at the board. However, the rocks are too heavy to move. The mere fact that it is *possible* that, say, a small landslide could happen to shift one of these rocks into exactly the right position, enabling Ella to exercise her dart-throwing competence, is surely not sufficient to make it feasible for her to hit the bullseye. (If that is a possibility, then there is a possible scenario in which Ella hits the bullseye in a way creditable to her competence: the shot would still manifest her darts-throwing competence, despite only being made possible by a freak accident.) The possible successful bringing about of an outcome (hitting the bullseye) is too modally fragile; it is not *safe*. (Thus, in the above snooker case, it is *not* feasible for me to hit the shot in one attempt, even though I could, of course, get lucky and hit it first time. But it *is* feasible for me to hit it given a number of attempts.)

---

[8] I presume one can be competent to bring about an outcome that one cannot bring about *intentionally*, as in the cases discussed above.

[9] It may be, though I am not sure, that the notions of competence and safety I make use of here have to be spelled out in terms of probability in some way. This would not turn my account into a probability-based account. If probability does come into my account, what matters is that it be *possible* for an agent to act in a way that gives the outcome a certain probability, not that it be probable *tout court* or probable conditional on trying.

[10] Note that the requirement is that *the agent* bring about O safely, not that the action do so.

There is some overlap in the work that can be done by the safety and competence requirements. A simpler account that only introduced one or other of these notions would suffice to get things right in most cases. Still, though, both are necessary. The need for the competence requirement, in addition to the safety requirement, is brought out by a case of Sosa's:

> A protecting angel with a wind machine might ensure that [an] archer's shot would hit the bullseye … and a particular shot might hit the bullseye through a gust from the angel's machine, which compensates for a natural gust that initially diverts the arrow. (Sosa 2007, p. 29)

The archer's shot hits the bullseye *safely* in this case; thanks to the 'protecting angel', it is not the case that the archer could easily have failed to hit the bullseye. But we do not want to say that it is *feasible* for the archer to hit the bullseye given the natural gust of wind diverting her arrows, simply because it is *possible* that a 'protecting angel' could intervene in this way (and so, it is possible that the archer hits the bullseye safely). It needs also to be possible for the archer to hit the bullseye in a way that manifests her competence. And it is not, if the only possibility in which she succeeds is one that involves the protecting angel. Thus, an outcome is only feasible for an agent if it is possible for them to bring it about *both* safely *and* competently.

Conclusion

There are many ways, then, in which the concept of feasibility can be specified, but these specifications are differentiated by a single variable, the facts held fixed. The facts we choose to hold fixed could, in principle, be any set of facts of the world, but, of course, some such sets of facts will be relevant for certain sorts of normative theory, while others will not. If we are interested, for instance, in the feasibility of bringing about a system of participatory democracy in the UK in the next decade, there are various questions we could have in mind. At quite a rough level of characterisation, we might hold fixed, for instance, the existing parliamentary system, the voting system, the configuration of political parties and the state of technology (as well as, no doubt, some more basic facts about physics, biology and so on). Then, to answer *this* feasibility question, we have to assess whether it is possible, *compatibly with these facts*, for us to bring about the desired system within a decade safely and competently.

The principal claim of this paper, though, was that the concept of feasibility should be understood as multivocal. The data on use of the concept in ordinary language and moral reasoning, I argued, point in this direction, and available attempts to give a *univocal* account encounter problems. In the absence of a univocal account that can successfully explain away the appearance of multivocality, we have good preliminary reason to suppose that the concept is multivocal.

What *are* the implications, then, of a multivocal account for moral and political philosophy? No single one of the many possible specifications of the concept is obviously privileged as the one relevant to practical deliberation. Which specification is meant by a particular feasibility claim may be determined by the context.

Alternatively, it may be indeterminate, requiring further specification for the claim to have determinate truth conditions. If this is correct, then to reject a moral or political theory (as either incorrect or useless) it will not be sufficient simply to say that its realisation is not feasible. That is true independently of any view one might have about theory being independent of feasibility facts. To say, for instance, merely that participatory democracy is not feasible is not to say anything determinate, until we have specified (or have tacitly understood) a specification of what we mean by 'feasible'. In the vast majority of cases, outcomes will be feasible on certain specifications (holding fixed certain sets of facts) and not on others. Thus, we cannot reject proposals simply by saying that they are infeasible. Rather, we need to know holding fixed which facts they are infeasible (i.e. given which FCs), and we need an idea of according to exactly which specifications a proposal needs to be feasible in order to be acceptable for a particular kind of moral or political inquiry.

Compliance with Ethical Standards

References

Brennan, Geoffrey and Southwood, Nicholas (2007) 'Feasibility in Action and Attitude' in T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson (eds) *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz.* http://www.fil.lu.se/hommageawlodek.

Erman, Eva, and Niklas Möller. 2020. A world of possibilities: The place of feasibility in political theory. *Res Publica* 26: 1–23.

Estlund, David. 2011. Human nature and the limits (if any) of political philosophy. *Philosophy and Public Affairs* 39: 207–237.

Gilabert, Pablo. 2009. The feasibility of basic socioeconomic human rights: A conceptual exploration. *The Philosophical Quarterly* 59: 659–681.

Gilabert, Pablo. 2017. Justice and Feasibility: A Dynamic Approach. In *Political Utopias: Contemporary Debates*, K. Vallier and M. Weber (eds). Oxford: Oxford University Press.

Gilabert, Pablo, and Holly Lawford-Smith. 2012. Political feasibility: A conceptual exploration. *Political Studies* 60: 809–825.

Ichikawa, Jonathan Jenkins and Steup, Matthias (2017) 'The Analysis of Knowledge' in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2017 edition), URL = <https://plato.stanford.edu/archives/spr2017/entries/knowledge-analysis/>.

Jensen, Mark. 2009. The limits of practical possibility. *The Journal of Political Philosophy* 17: 168–184.

Lawford-Smith, Holly. 2012. The feasibility of collectives 'action'. *Australasian Journal of Philosophy* 90: 453–467.

Lawford-Smith, Holly. 2013a. Understanding political feasibility. *The Journal of Political Philosophy* 21: 243–259.

Lawford-Smith, Holly. 2013b. Non-ideal accessibility. *Ethical Theory and Moral Practice* 16: 653–669.

Lewis, David. 1973. Counterfactuals and comparative possibility. *Journal of Philosophical Logic* 2: 418–446.

Pritchard, Duncan. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy* 109: 247–279.

Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, vol. I. Oxford: Oxford University Press.

Southwood, Nicholas. 2016. Does 'ought' imply 'feasible'? *Philosophy and Public Affairs* 44: 7–45.

Stemplowska, Zofia. 2016. Feasibility: Individual and collective. *Social Philosophy and Policy* 33: 273–291.

Wiens, David. 2015. Political ideals and the feasibility frontier. *Economics and Philosophy* 31: 447–477.