Canadian Journal of Philosophy (2025), 1–15
doi:10.1017/can.2025.10013

ARTICLE

A Multifunctional Account of Political Feasibility

Brian Carey

School of Government and International Affairs, Durham University, Durham, UK
Email: brian.carey@durham.ac.uk

Abstract
In this article, I argue for a "multifunctional account" of political feasibility and against recent attempts by several theorists to defend accounts of political feasibility that reduce feasibility judgments to a single function. According to the view I defend, political feasibility can (and should) serve multiple useful functions in our political deliberative practices. This pluralist and context-dependent approach allows us to retain the insights provided by various monofunctional accounts, while avoiding the limitations of each.

Keywords: feasibility; nonideal theory; political strategy; public deliberation; pluralism

1. Introduction

This article is about what we usually mean (and what we ought to mean) when we say that some particular political proposal is "feasible," or "infeasible" or "more feasible" or "less feasible" than some alternative. Claims about feasibility (whether explicit or implicit) are a common feature in public political debates, though the normative significance of a proposal's feasibility is often ambiguous or under-specified. For example, feasibility claims are likely to be invoked when considering political proposals that aim at radical (especially unprecedented) institutional change, perhaps in an effort to argue that such change is impossible because it conflicts with unalterable facts about human nature, or to reassure the audience to the contrary.[1] In these cases, feasibility is treated as a binary concept that demarcates the boundary between proposals that are possible and those that are impossible, in service of the underlying normative principle that "ought" implies "can". To declare a political proposal to be infeasible in this sense is to declare that there is no moral obligation to pursue it, while to declare it as feasible leaves open the possibility that such an obligation might exist.[2] In other cases, feasibility claims are invoked when considering which of a range of possible options is to be pursued, on the basis that some options are more likely to be successful than others.[3] Such claims may involve practical questions of resource allocation, but these also contain a normative dimension insofar as it matters morally that we use limited resources effectively and efficiently. Such claims may also implicate normative principles concerning the

[1]Examples of such proposals that have been offered in the literature include universal basic income (Southwood, 2022, p. 121), the eradication of poverty, combating global climate change (Stemplowska, 2021, p. 2399), communal models of child-raising, and "pretax max" theories of justice (Estlund, 2011, pp. 211–215).

[2]In fact, there may be circumstances in which we are obliged to *pursue* an impossible goal even if we are never obliged to *achieve* it, if the pursuit itself leads to desirable outcomes. However, public debates about political feasibility tend to assume implicitly that this is not the case for the proposals in question.

[3]For more on the differences between binary and scalar concepts of political feasibility see Gilabert and Lawford-Smith (2012).

allocation of risk—strategies that prefer more feasible proposals over less feasible proposals will tend to minimize risk, though this is often complicated in cases where we consider a proposal that is more feasible but leads to a less desirable outcome than another proposal that is less feasible but more desirable. In some cases, feasibility claims are invoked not as a way of arguing for or against the merits of a political proposal, but as a means of explaining why some particular agent is (or would be) morally accountable for failing to successfully implement the proposal in question. Public political debates on each of these matters are further complicated by the epistemology of feasibility—there is of10 deep (and often reasonable) disagreement about whether a particular proposal is feasible and/or how feasible it is, and what we ought to do in the face of uncertainty about the reliability of our feasibility judgments independently of the normative implications that follow from feasibility judgements themselves.[4] In sum, questions about political feasibility are ubiquitous, ambiguous, and important.

Philosophers have responded to these complexities by offering conceptual theories of political feasibility that aim to clarify the concept in a way that aligns with our linguistic practices and moral intuitions. According to most of these approaches, a good conceptual theory of political feasibility is one that is compatible with how we tend to talk about feasibility[5] and helps to explain the normative significance of feasibility considerations. This has led to what we might call the "*feasibility as…*" debate, with plausible candidate theories including feasibility as probability of success conditional on trying (Estlund, 2008, 2011), all-purpose resource availability (Wiens, 2015), likelihood of success given the availability of incentives (Stemplowska, 2021) and deliberation-worthiness (Southwood, 2022).

In this article, my aim is to argue for the mutual compatibility of these approaches by identifying and arguing against an assumption that is common to each: that political feasibility judgements serve (and ought to serve) a single fundamental function in the context of public political deliberation.[6] Instead, I will argue for a "multifunctional" approach to political feasibility, which holds that feasibility judgments can (and ought to) serve several different purposes, depending on the context. This account yields a fundamentally pluralist understanding of political feasibility, allowing us to preserve the insights that have been developed via monofunctional accounts while avoiding the limitations of each.

I begin in Section 2 by examining a recent theory of feasibility offered by Nicholas Southwood, which is notable because of his explicitly functionalist approach (which I will endorse) and for his suggestion that feasibility judgments are ultimately about "deliberation-worthiness" (which I will reject in favour of a pluralist view that includes but is not limited to this function). In contrast to Southwood, prominent alternative approaches to feasibility have tended to focus on the modal properties that determine whether and to what extent a particular state of affairs should count as feasible for some particular agent(s).[7] Nevertheless, I will argue that these alternative approaches reveal alternative functions that feasibility judgments play in our political deliberative practices. In Section 3, I consider theories of feasibility offered by David Estlund and Zofia Stemplowska and argue that these suggest feasibility judgments can serve an important blame-allocating function and that the gaps in these approaches reveal a further function for feasibility judgments as predictions. In Section 4, I turn to David Wiens' account of feasibility and argue that it suggests a further comparative function that feasibility judgments sometimes play in our political deliberations. While this analysis does not necessarily represent an exhaustive account of feasibility's functions,

[4]See, for example, Carey (2015).

[5]However, Southwood notably dissents from approaches he refers to as "beholden to practices of linguistic attribution" (Southwood, 2022, p. 161).

[6]Southwood (2022) distinguishes his explicitly "functionalist" approach from those of his rivals—however, I will suggest below that alternative functionalist approaches can be extracted from these rival views (though the functions their views imply for feasibility judgments differ from that identified by Southwood).

[7]I am grateful to an anonymous referee for helping me to clarify this point.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

identifying the importance of such functions should be enough to show that we should be both functionalists and pluralists when it comes to political feasibility. In Section 5, I explain what this approach might look like in practice and consider potential objections that a multifunctional approach might be less clear than a monofunctional alternative and that it may fail to provide us with enough guidance when it comes to choosing specific conceptions of feasibility for specific functions. Section 6 concludes.

2. Functional Feasibility and Feasibility as Deliberation-Worthiness

In the introduction of this article, I described various ways in which we appeal to feasibility considerations in the context of public political deliberation. To think about feasibility in this way is to begin to take a functionalist approach to the concept—to conceive of feasibility claims (or judgments) as something we deploy when we engage in political deliberation.

Nicholas Southwood argues for a functionalist approach to feasibility because he believes it offers a better prospect for agreement than what he calls the "linguistic approach." The latter tries to align a philosophically robust account of feasibility with our typical linguistic practices, while the functionalist approach instead focuses on our "practices of practical deployment" (Southwood 2022, p. 123). While I hope Southwood is correct on this point, my own reason for preferring a functionalist approach is grounded in what I see as its action-guiding potential for political agents. The functionalist perspective requires us to attend first and foremost to political practices, and to conceive of the idea of political feasibility as a kind of tool that is deployed in deliberative contexts in order to achieve certain aims. Identifying these aims—as the functionalist approach requires us to do—allows us to explain what feasibility claims can do and how they can do it. A good functionalist approach to political feasibility has the potential to act as a kind of instruction manual for political actors, by clarifying the way(s) that feasibility claims may be deployed (albeit such guidance will be incomplete without a substantive normative element).

Although Southwood's view is my primary target here, he is not the first to argue for functionalism about feasibility - Eva Ermann and Niklas Moller's work is described by Southwood as "the best-developed articulation of a functionalist approach [to feasibility] to date" although he considers it to be "programmatic and promissory" (Southwood, 2022, p. 125). Ermann and Moller argue for two broad metatheoretical constraints on feasibility theorizing, the first of which is a "functional constraint" which is "the requirement that the guiding principles of a normative account must be appropriate for what the account aims to do" (Ermann & Moller, 2020, p. 8).[8] In their discussion of appropriateness, they distinguish between types of normative principle (such as a principle of justice, democracy, or political legitimacy), domains of application (such as a family, a state, or the world) and distinguish between theories that aim to provide guidance in the present and those that are more future-oriented. The view that I will develop in this article aligns with Ermann and Moller's in its commitment to a functionalist approach that considers feasibility judgments to be highly flexible and applicable in a wide range of contexts (in contrast to Southwood's monofunctional account). My approach will differ from theirs in that they are primarily concerned with establishing the flexibility and context-sensitivity of feasibility judgments in general, whereas I am interested in identifying and examining specific functions in addition to the deliberation-determining function identified by Southwood. While I take it that nothing I say here is fundamentally incompatible with Ermann and Moller's position, the examples they discuss do not explicitly address feasibility's predictive, blame-allocative, or comparative functions as I shall discuss below, nor feasibility in its deliberation-worthiness determining function, as identified by Southwood. My aim is to build on the functionalist foundations established by Ermann and Moller

[8]The second constraint is a "fitness constraint" which tracks consistency between one's different normative commitments (Ermann & Moller, 2020, pp. 8–9).

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

(and developed by Southwood) by providing a clearer map of the territory covered by feasibility judgments.

Southwood's functionalist account conceives of feasibility as "deliberation-worthiness" such that to judge a political proposal as feasible is to judge it worthy of consideration (because it is pointless or irrational to spend time considering whether we should do something we believe to be infeasible). As a description of our deployment practices, this might suffice insofar as we limit it to an account of feasibility in its binary sense, where feasibility and infeasibility map onto the concepts of possibility and impossibility. Certainly, when we deliberate about which political proposal to pursue we generally assume that all of the options on the table are actually possible, even if some are less likely to be successful than others. However, once we consider the role of scalar feasibility judgments (which track probability rather than mere possibility) it seems clear that these serve many functions beyond determining deliberation-worthiness. For example, when we are considering a range of possible options we use scalar feasibility considerations to help us decide among them—thus feasibility considerations play an important role throughout our deliberations, and not merely in deciding initially which options are worthy objects of deliberation.[9]

If this is so, then it is difficult to see why we should count deliberation-worthiness as the sole defining function of feasibility assessments to the exclusion of the alternatives I will discuss below. It will not do to suggest in response that the deliberation-worthiness function is in a sense prior to other feasibility functions insofar as these may involve judgements about possible courses of action that have already been deemed worthy of deliberation. That may be so, but it does not follow from the fact that one function of feasibility judgments presupposes another that feasibility judgments have only one function.

What I am arguing for in this article is, in effect, what Southwood (2022, p.160) identifies as "practice-based pluralism" about feasibility, according to which there are multiple roles that feasibility judgments can play, none of which are reducible to any of the others. In an effort to reject such views, Southwood considers the possibility of a prescriptive form of feasibility whose function is to serve as advice for what one ought to do and a hypological form whose function is to determine whether one ought to be blamed on the basis of the feasibility (or infeasibility) of one's actions (or inactions). In the next section, I will consider the hypological form in more detail in the context of Estlund and Stemplowska's views about the relationship between feasibility and blame-allocation, but for now, here is what Southwood says about the idea that one function of feasibility judgments is to track liability to moral criticism:

> We frequently criticize others for failing to do what we take to have been infeasible for them to do at the time: say, if they intentionally or negligently made it the case that it was not feasible for them to do it; or if they intentionally or negligently failed to make it the case that it was feasible for them to do it; or if it was infeasible for them to do it because of certain criticizable character traits. (Southwood, 2022, p. 161).

Southwood's aim here is to show that our judgments about liability to moral criticism work independently of our judgments about feasibility (and thereby to rule out the idea that feasibility serves a blame-allocating function). I am not convinced by these examples. The first variations involve an important but underexamined element of temporality, where we are inclined to blame a person now because of something they did (or failed to do) in the past. It seems to me that the reason

[9]Indeed, binary feasibility judgments alone are not always decisive when it comes to determining deliberation-worthiness—possible but highly unlikely plans of action (such as winning the lottery to raise funds for a campaign) may be regarded as unworthy of deliberation because they are insufficiently feasible in the scalar sense, even if they are feasible in the binary sense. As Southwood notes, for feasibility judgments to provide deliberation-guiding recommendations "it must be the case that "feasible" implies "not counterfactually fluky"" (Southwood, 2022, p. 141).

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

we are inclined to blame them now is precisely because there was some point in the past at which (we assume) it was feasible for them to act in the appropriate way. If in fact it was never feasible for the person to act in the past, then I suggest we would not be likely to blame them in the present. The second example involves infeasibility that is caused by the presence of "certain criticizable character traits." What might these be, if not traits that it is (or was) feasible for the person to change at some point in time? While feasibility is certainly not sufficient for moral blame (there might be other factors that tell against blame when all things are considered), it does seem to be a necessary ingredient: it seems to me that the person in these examples cannot be morally liable unless feasibility enters the picture at some point in the story, which suggests that feasibility does have an important hypological (or blame-allocative) function. I return to the idea of feasibility as blame-allocation in Section 3.

Here is what Southwood has to say about the possibility of a "prescriptive" function for feasibility:

> Similarly, it would seem that we might very well advise (or enjoin or exhort) someone to do something that we take to be infeasible without this striking us as remotely odd or inappropriate—so long as we were to supplement our (unconditional) advice about what to do with additional conditional advice about what to do insofar as they do not follow our (unconditional) advice (Southwood, 2022, p. 161).

Here Southwood aims to show that feasibility does not play a prescriptive role by showing that we sometimes advise people to do that which we take to be infeasible. I am more sympathetic to this point, since we do not think that claims such as "φ is feasible" implies "you ought to φ." Obviously there is no such implication when φ represents, for example, a trivial or morally contemptible act, but even when it represents some very positive act it does not follow (there might be even better things you should attempt even though they are less feasible, or φ might be supererogatory, for instance).

A modified form of this prescriptive function may fare better, however. Suppose we were to consider "φ is feasible" to imply something like "all else being equal, you have a reason to choose φ over less feasible alternatives." Feasibility judgments do seem to serve this function in cases where the rest of the moral equation has already been "filled in," so to speak. In other words, there are cases where we are trying to determine which route to follow to arrive at some particularly morally valuable destination for which feasibility judgments serve as a kind of tiebreaker, or at least serve to enhance the value of a proposal that aims to realize some morally desirable state of affairs. In these (admittedly restricted set of) cases, feasibility judgments can serve a prescriptive role, at least insofar as they add weight to the scales of our moral judgements.

In the following sections, I will discuss alternative functions for feasibility judgments in more detail, but there are some general reasons to prefer a pluralist account over a monist one that are worth noting at this point. First, pluralism about feasibility's functions offers a more flexible approach to public deliberation, allowing us to tailor the relevant function to the task at hand rather than trying to adapt a single function for multiple deliberative contexts. Second, pluralism better reflects our existing practices- as we have already seen, we use feasibility judgments for other tasks besides determining deliberation-worthiness, so a pluralist approach does a better job of capturing our actual deployment practices, which is a key aim of a functionalist approach. Third, the pluralist approach I argue for in this paper is non-exhaustive in that it does not exclude the possibility of identifying additional functions beyond those discussed here, making it more amenable to revision in the future.

To summarize, I think Southwood is right to take a functionalist approach to feasibility, but wrong to think that feasibility is only about deliberation-worthiness. I am convinced by his suggestion that feasibility judgments function to determine deliberation-worthiness but not by his attempts to rule out other roles for feasibility, such as blame-allocation or (in at least some cases)

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

prescription. In the next section, I will examine in more detail the idea that feasibility judgments can serve an important role when it comes to allocating blame.

3. Feasibility as Blame-allocation and Feasibility as Prediction

> It is pretty easy to dance like a chicken in front of your boss. Put your hands up under your arms, thrust your head forward rhythmically, and so on. It is easy, but you and I both know you will almost certainly never do it. The same goes for certain things that might be morally required. (Estlund, 2008, pp. 13–14).

Cases such as Estlund's are often taken to provide a decisive objection to "simple probability-based accounts" (Southwood, 2022, p. 127) of political feasibility, according to which something counts as feasible for some particular agent to the extent that it is likely that they will bring it about.[10] As Estlund's example shows, there is a distinction between how easy it is for us to do something if we attempt it and how likely it is that we will actually attempt it.

While there are several alternatives to the simple probability-based account[11], their advocates appear united on the reason why such an account must be rejected: it lets people off the hook in cases where they fail to perform some morally important action merely because they will not muster the will to attempt it. These are cases of *motivational failure*, and (provided that the reason for the failure is not pathological[12]), opponents of simple probability-based accounts believe that motivational failure does not exculpate.[13] Such theorists accept some form of the principle that "ought implies feasible": we are not morally obliged to do the impossible, or the highly improbable (though sometimes we may be obliged to try).[14] On this view, to show that some action is infeasible for you is equivalent to showing that it is impossible for you and that you could not be obliged to perform it, or (in the case of the highly improbable) that you could not be blamed for failing to achieve if you try. If non-pathological motivational failure does not exculpate, then the correct theory of feasibility must not count motivational constraints as constraints on feasibility lest we render the wrong judgments about agents' moral blameworthiness in cases where the only thing stopping from acting is their own lack of will.

Estlund's preferred approach is to move to a *conditional probability* account of feasibility, according to which an action is feasible for someone if they would tend to succeed were they to try and not give up. This explicitly detaches feasibility judgements from questions of motivational failure, though Estlund does add the caveat that this does not include rare and extreme "pathological" cases where he thinks that motivational failure can be exculpatory. It also provides us with the intuitively correct answer in cases involving flukes: while it may be possible for me to hit a bullseye on a dartboard while blindfolded, the sheer improbability of me doing so (were I to try and not give up) suggests that it would be infeasible for me to do so, and that it would be improper to blame me for failing to succeed.

Similarly, Zofia Stemplowska has argued for an *incentives* account of feasibility, according to which an action is feasible for some agent "if there could be an incentive such that it would

[10]See, for example, Brennan and Southwood (2007, pp. 8–10), Gilabert and Lawford-Smith (2012, p. 824), Wiens (2015, p. 451), Stemplowska (2016, p. 274), and Southwood (2016, p. 11).

[11]Southwood (2022) provides a helpful overview of some of the most plausible possibilities, including Estlund's and Stemplowska's models discussed in this article.

[12]For example, Estlund (2011, pp. 230–231) takes a very strict view with regard to the non-exculpatory nature of motivational constraints, but is still willing to allow that pathological motivational inabilities can exculpate.

[13]For a contrary view to Estlund's, see Carey (2016).

[14]As Southwood puts it, "[f]or any agent A, action X, and time t, A ought to X at t only if it is feasible for A to X at t." (Southwood (2022, p. 126). Southwood (2016) distinguishes between different kinds of "ought", only some of which need imply feasibility, but it is the feasibility-implying oughts that concern us here.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

incentivize the agent to try to perform the action and the agent is likely to succeed," though she notes that "[i]t must be the case that the incentive is conceptually possible but not that it is available to be offered" (Stemplowska, 2021, p. 2390). This, Stemplowska argues, allows us to better distinguish between cases of genuinely pathological motivational failure (for which a person should not be blamed) and "mere unwillingness" for which a person might be blamed.

On both accounts, the aim is to provide a theory of feasibility that aligns with our intuitions about blame (and in particular, that a mere lack of motivation to act is not exculpatory). While neither frame their approach in explicitly functionalist terms, Estlund and Stemplowska's concern that their theories align with our intuitions about blame reflects the fact that we often use claims about feasibility to help us determine blameworthiness.[15] We can thus derive functionalist versions of their accounts by considering which modal properties they take to be relevant and why—in other words, once we have an account of the relevant modal properties, what do we want to do with that information in the context of political deliberation? For Estlund and Stemplowska, it seems the answer is that information about feasibility is to be used to help us determine our moral obligations and to identify those who deserve blame for failing in their moral duties (while avoiding blame for those who can appeal to genuinely exculpatory reasons).

At this point, we are presented with two monofunctional accounts of feasibility, which each identify different functions for feasibility judgments—determining deliberation-worthiness, on Southwood's account, and allocating blame on Estlund's and Stemplowska's accounts. A multi-functional approach to feasibility suggests that each approach provides strong reasons to reject the monofunctional character of the other, but not that these accounts are exhaustive when it comes to identifying the functions of feasibility judgments. In the next section, I will suggest that feasibility judgments can also function as predictions or comparisons in ways that are just as useful as their deliberation-worthiness-determining or blame-allocative functions. To appreciate the usefulness of these further functions, however, it will be helpful to first explore some of the limitations of monofunctional accounts of feasibility as blame-allocation besides the fact that they neglect the deliberation-worthiness determining function as identified by Southwood.

First, we may note that feasibility judgements involving blame-allocation are often made "after the fact," that is, once some particular course of action has failed and we engage in a kind of moral autopsy to determine who is to blame and to what extent. This presents a contrast with the deliberation-worthiness determining function identified by Southwood's approach, which suggests that feasibility judgments typically precede deliberation (though recall that a limitation of Southwood's approach is that feasibility judgments apply only up to this point and do not help us to determine which deliberation-worthy course of action to pursue). Feasibility judgements involving the allocation of blame are more temporally and counterfactually flexible in the sense that we can use them to say who would be blameworthy if an action were to fail, or who is blameworthy once an action has failed. However, such theories also lack predictive power as a result: Estlund's concept of feasibility predicts what would happen if the relevant motivations are present (they may or may not be in practice) while Stemplowska's predicts what would happen if the relevant incentives were offered (they may or may not be available to be offered in practice). This is not a problem if our aim is to determine the extent to which a person is blameworthy for failing to bring about some particular state of affairs, but this in turn suggests another important role for feasibility judgments —we sometimes use feasibility claims to make predictions about what is likely to happen, and to plan our political strategies accordingly. Both Estlund and Stemplowska's accounts demonstrate the importance of avoiding feasibility as actual probability if we are interested in the allocation of blame,

[15]In her critique of Estlund's approach, Stemplowska says that her approach "preserves the spirit of [Estlund's] account in that it does not get people off the hook with regard to whether their acting is feasible merely if they are unwilling to so act…" (Stemplowska, 2016, p. 291). I take this as evidence that a focus on blame-allocation is central to both Estlund and Stemplowska's accounts.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

but in making this move feasibility judgments become incapable of serving an important action-guiding role based precisely on the actual probability considerations we must set aside for the allocation of blame (nor is this a role that Southwood's feasibility judgments can play, once we have determined which options are worthy of deliberation).

Consider how often we treat feasibility judgments as reasons to choose one political proposal over another. When all else is equal, it is natural to say that one proposal being more feasible than the alternatives gives us a decisive reason to pursue it. This can make sense only if we suppose that our feasibility judgments (when used in this way) are tracking our judgments about actual probability rather than conditional probability (or possible incentives). Such judgments take into account not just probability of success conditional on trying but the actual probability of trying itself. When applied in this way, feasibility judgments concern the fair and efficient allocation of resources, given a range of options, some of which are more likely to succeed than others.[16]

Both Estlund and Stemplowska's views fail to capture this action-guiding function of feasibility judgments because they allow for a course of action to count as feasible for an agent even when there is zero chance that the agent will actually attempt to pursue it. Again, this is not a problem if our aim is to allocate blame—we can already know everything we need to allocate blame at the point where we have determined probability of success conditional on trying. However, if our aim is to determine whether to pursue a particular political proposal, it is not enough to know that it is likely to succeed if we try—we also need to know what the probability is that we will actually try. If we stick with Estlund or Stemplowska's conceptions and confine feasibility judgments to a single blame-allocating function, we may find ourselves in a position where we are able to determine that some valuable course of action is highly feasible, but not whether we have any reason to pursue it. This is at odds with our practice of using feasibility judgments in ways that suggest that to determine a proposal to be highly feasible is to provide a reason to prefer it over similarly desirable but less feasible alternatives. In other words, we often treat feasibility judgments as though they give us reasons to act, to prefer one proposal over another, and so on, but Estlund and Stemplowska's conceptions limit feasibility such that feasibility judgments provide us with reasons to blame, but not reasons to act in the absence of further information about the actual probability of success.

One might worry that there is something incoherent in the idea of applying predictive feasibility judgments to our own conduct in the context of practical deliberation.[17] Motivation tends to be something that leads to deliberation, or follows from deliberation, but it is not usually an object of deliberation itself. Nevertheless, I suggest that there are at least some cases where it can be useful to make predictive feasibility judgments about our own motivations. The first involves cases where we predict that our motivations may be unstable over time—for example, suppose you believe that you tend to change your mind a lot, or that you are particularly weak-willed such that you often find it difficult to maintain commitment to demanding goals, or any number of other features that lead you to believe your motivations may not hold in the long run. Similarly, a second type of case involves judgments about the reliability of our own ability to predict motivational failure. For example, suppose that I am considering a course of action that will require me to sacrifice more than I have ever sacrificed before and I am uncertain whether I will be able to manage this. Depending on the context, these kinds of predictive judgments may have a strong bearing on our chosen course of action (most obviously—they may give us reason to take measures to reinforce our commitments in the future or to make it more difficult for us to waver from them).[18]

The conception of feasibility that is implicated in strategic judgments about resource allocation is thus sensitive only to the actual motivations (or lack thereof) of those involved, in contrast to

[16]See Wiens (2015).

[17]I thank an anonymous reviewer for pressing me on this point.

[18]I suspect these kinds of strategies will be particularly familiar to those of us who have tried and failed multiple times to stick with diet and exercise programs by (for example) publicly announcing our intentions to our friends or conscripting a gym buddy in order to make it more difficult for us to renege on our commitments.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

Estlund's conception, which determines feasibility independently of actual motivations. Nor does this conception involve questions about whether there are any possible incentives that might be provided to motivate agents to act but rather whether there are any actual incentives that are available. Where there is a pro tanto duty to maximize the chance of success, this requires agents to attend to the actual probabilities of success given the best information available, not counterfactual probabilities that anticipate a lack of motivational constraints or the presence of hypothetical incentives.[19] Political actors may regret such motivational failures when they occur, they may criticize others (or themselves) for their weakness of will, or try to change their motivations, but their strategic decisions must nonetheless remain constrained by probabilistic predictions about their chances of success, given the relevant facts (including facts about people's actual motivations). In this sense, "predictive feasibility" understood as probability of success is narrower in scope, and more limited in its application than conceptions based on counterfactual probabilities or conceptually possible incentives.

Suppose, for example, that a group of activists is working to abolish capital punishment (and that this would be a highly morally valuable outcome). Suppose furthermore that they have good reason to believe that appealing to the inherent dignity of all human beings is unlikely to motivate a sufficient number of voters, while appealing to the amount of money that is spent administering capital punishment is likely to be far more persuasive.[20] It would be irresponsible of the campaigners to discount these facts about voters' motivational failures in their strategic deliberations, and to act as if they made no difference to the feasibility of the options before them.[21] This is consistent with holding that voters could bring themselves to do otherwise, and with blaming them for their failure to do so, because feasibility judgments in this context are playing a purely predictive role, which tracks only whether a particular course of action is likely to succeed, but not why. This is not to say that there is any inherent tension between the predictive function of feasibility and the blame-allocative function: we can simply hold that in cases such as this one it would be entirely feasible for the voters to vote for (what we can stipulate to be) the right reasons, but not feasible for the campaigners to make them do so. Motivational failures may not count as constraints on feasibility for those who experience them (for the purpose of blame-allocation), while counting as constraints on feasibility for the rest of us (for the purpose of guiding our political strategies), and this is not at all surprising. Rather, the point is that when we wish our feasibility judgments to serve a predictive function, motivational failures have no special status among feasibility constraints, in contrast to feasibility judgments that serve a blame-allocative role, where motivational failures may be uniquely non-exculpatory compared to other kinds of constraints.

Again, I emphasize that none of these points are intended to serve as criticisms of blame-allocative models like Estlund's or Stemplowska's, but rather to show that there is a domain of decision-making where we simply do not need to rely on these models to render useful, accurate, and morally significant (predictive) judgments of feasibility. Blame-allocative models undoubtedly serve an important function, but the move from a simple probability-based account of feasibility to these alternatives is necessitated only to the extent that we are interested in allocating blame. When it comes to making predictive feasibility assessments, we must attend to actual and not counterfactual probability.

[19]Of course, *some* consideration of incentives will be strategically important, but Stemplowska's view is not restricted to just those incentives that agents might be able to provide, but includes any incentive that could conceivably be provided, whether it is in the power of the relevant agents to provide it or not.

[20]I emphasize that these are stipulations for the sake of argument and not empirical claims about what kinds of approaches real voters are likely to find most persuasive.

[21]To be clear, the activists need not treat these predictive feasibility judgments as generating decisive reasons all-things-considered (it may be that despite one strategy being much more likely to succeed, there are other good reasons not to pursue it). My claim here is just that such predictions must be taken into account as part of the activists' deliberations.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

We can see how the predictive and blame-allocative approaches diverge in Estlund's original case of the unlikely chicken-dancer: the blame-allocative model holds that it is feasible for you to dance like a chicken, to the extent that you would succeed were you to try[22] while the predictive model holds that any course of action predicated on the assumption that you will dance like a chicken is likely to fail, given that you are highly unlikely to try. Each model diverges when it comes to the significance attached to the stipulation of motivational failure—the blame-allocative models regards this as irrelevant for feasibility because it is irrelevant for your potential blameworthiness should your failure to act constitute a serious moral wrong[23], while the predictive model regards your lack of motivation as a feasibility constraint on any plan of action that requires you to dance like a chicken. It is important to emphasize that this divergence does not lead to a contradiction between both kinds of feasibility judgment but rather to a difference in focus—the blame-allocative judgment focuses on facts indicating moral blame while the predictive judgment focuses on facts indicating probability of success.

To summarize, I agree with Estlund and Stemplowska that feasibility judgments help us to allocate blame, and that in doing so we should attend to counterfactual conditionals rather than actual probability (either probability of success conditional on trying or probability of success conditional on the provision of possible incentives). However—as with my analysis of Southwood's account - I reject the idea that this is the *only* function that feasibility judgments serve in our political deliberations. Where we aim to provide action-guiding recommendations for political actors, our feasibility judgments must attend to actual probabilities instead.

4. Feasibility as Comparison

Predictive feasibility judgments undoubtedly have their limitations too. As we have seen in the previous section, predictive judgements tell us nothing about how to allocate blame among those who fail to achieve or attempt what we think they ought to have achieved or attempted—this is a useful function that approaches such as Estlund's and Stemplowska's can serve instead. Even more obviously, perhaps, the mere fact that some political proposal has a high probability of success is not in itself a reason to pursue it, but rather only serves to add weight to whatever value is already inherent in the proposal itself.[24] Where a proposal is both feasible and desirable, it may sometimes be rejected in favour of another that is more desirable, though less feasible (or vice versa), depending on the context.

Yet even on their own terms, predictive feasibility judgments are at best only one part of the action-guiding equation. Paradoxically, there is risk in risk-aversion: overly-pessimistic predictive feasibility judgments may guide us toward outcomes that are less valuable than we might otherwise have achieved and following these routes may create issues of path dependency that make more valuable alternatives more difficult or impossible to achieve in the future. Predictive feasibility judgments may also undermine the politics of the unprecedented, especially if we fail to take proper account of the limitations on our ability to make accurate feasibility predictions.

One area where these dangers are particularly clear involves methodological questions concerning ideal and nonideal theory in contemporary theories of justice.[25] Pablo Gilabert and Holly Lawford-Smith (2012) distinguish between "hard" and "soft" feasibility constraints on a theory of justice, where the former represent assumptions about our theory that we assume to be permanent, unalterable features of the world, while the latter represent obstacles that can (in principle) be removed as we progress from the status quo to a more ideal state of affairs. Crucially, predictive

[22]In other words, we can dance if we want to (Doroschuk, 1982).

[23]I leave it to the reader to imagine the strange circumstances under which failing to dance like a chicken in front of one's boss might constitute a serious moral wrong.

[24]I have in mind here what Stemplowska calls "action-guiding recommendations" (Stemplowska, 2008, p. 324).

[25]For a general overview of these debates, see Valentini (2012).

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

feasibility judgments that fail to properly track this distinction (and/or fail to adequately assess the "softness" of soft constraints) risk providing us with predictive feasibility judgments that are too pessimistic on the one hand, or overly utopian on the other.

Most significantly for our present purposes, binary feasibility judgments that aim to distinguish the possible from the impossible are sometimes assumed to play a role not only in identifying the best possible state of affairs (that is one in which only hard constraints remain) but in guiding us toward them—an approach that David Wiens (2015, p.448) describes as a "target view" of normative political theory. As Wiens argues, however, the mere fact that we can describe an ideal state of affairs does not entail that it should be a target for political reform, if it falls beyond what we are capable of given the resources at our disposal, which establishes a feasibility frontier.[26] Rejecting target views, Wiens argues for an approach to normative political theory that focuses not the idea of progress toward an ideal of justice, but rather on a "backward-looking" idea of progress from injustice:

> Reoriented toward the analysis of failures, normative political philosophy undertakes the following clinical tasks: identifying states of affairs that are suboptimal with respect to some evaluative criteria; diagnosing the causes of the suboptimal states; evaluating potential remedies for circumventing or mitigating the identified failure, assessing the moral costs and benefits of alternatives; prescribing remedies that are (as far as we can tell) likely to leave open possibilities for future progress. (Wiens, 2015, pp. 471–472).

We need not decide between a forward or backward-looking approach to justice to note that Wiens' account reveals a further role for feasibility judgments that complements its predictive role, namely a comparative role between the political proposal in question, and superficially similar proposals. When we deploy feasibility judgments in this way, to say that a proposal is feasible is to say that it is sufficiently similar to some other endeavour that has been shown to be successful elsewhere. This comparative role is distinct from feasibility judgments as all-things-considered predictions, but it complements these predictions to the extent that comparative feasibility judgments establish precedents and to the extent that precedents provide reasons to predict success.

We are now in a position to appreciate both the usefulness and the limitations of predictive feasibility judgments. While predictive feasibility judgments are silent on matters of moral responsibility, they play a vital role in questions of resources distribution—by establishing the limits of what it is possible for us to do with our resources, by helping to identify the costs to our resources associated with particular plans of action, and (when combined with comparative feasibility judgments) by helping to identify relevant precedents, especially when aiming to rectify injustice.

5. Multifunctional Feasibility

At this point, we have seen feasibility judgments serve several different and important functions: at the very least, they help us to decide which political proposals are worthy of deliberation, to determine the correct distribution of blame in the event that feasible proposals fail, and to decide which feasible proposals represent the fairest and most efficient use of our resources. In each case, the function we choose leads us to attend to a specific set of facts that are relevant for serving that particular function but which may not be relevant or appropriate if deployed in service of other feasibility functions (if, for example, we were to focus on actual probabilities while trying to allocate blame, or counterfactual probabilities while trying to make predictions). In this section, I argue that

[26]Specifically, the "attainable total stock of all-purpose resources is a configuration of (institutional, cultural, financial) resources that could emerge from a transformation of our current resource stock by a series of conversions or investments that is available to us" (Wiens, 2015, p. 454).

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

we should embrace the plurality of feasibility functions, rather than try to capture each of these within a single concept of feasibility, or by trying to restrict the concept to just one function.

In short, there is no need to assume that the ways in which we deploy feasibility judgments in ordinary public political discourse need to be captured by a philosophically rigorous concept of feasibility that identifies feasibility judgments with only a single function. Provided that we are clear about what kind of judgment we are making, feasibility judgments can serve multiple, mutually compatible functions within the context of public political deliberation. We can deploy Estlund's conditional probability or Stemplowska's incentives account of feasibility to explain whether someone is deserving of blame in cases where they failed to act only because they lacked the motivation to do so. We can deploy a simple probability account of feasibility to explain why we ought to pursue some particular political strategy, on the basis that it is more likely to succeed than the available alternatives. We can deploy Wiens' account of feasibility to determine the range of options available to us given our resources, or to compare different courses of action to establish a precedent between a potential course of action and a previously successful one. We can deploy Southwood's account to help us determine which political proposals should be up for deliberation and which should not. This is not an exhaustive list—there may be other functions that feasibility judgments may play, and if so, the more the merrier. A multifunctional account allows us to pick and choose, depending on the context, which kind of judgment to deploy, and to retain the insights from all of these approaches in doing so. The only significant cost to this kind of pluralism is that we must be careful to be clear about what conception of political feasibility we are using, but this seems a comparatively small cost.

One advantage of this pluralist, multifunctional approach to feasibility is that it seems to align with our pre-theoretical linguistic practices, which deploy the term in all of the contexts identified above. As noted earlier, Southwood is sceptical of approaches that aim to capture our linguistic practices within a theory of feasibility, but the apparent failure of linguistic approaches may be precisely because these theories aim to capture pluralistic behaviour within a monist account of feasibility. Monist accounts correctly identify the ambiguity in our linguistic practices, but misdiagnose the problem, assuming that clarity requires a single account of "feasibility as…" that captures how we talk about, or deploy, feasibility judgments in our deliberative practices.

However, even if I am right that our ordinary linguistic practices reveal multiple roles for feasibility, it does not necessarily follow from this alone that we should continue to use feasibility in this way.[27] We might accept that feasibility is a multifunctional concept at present but believe that things would be clearer if we were to restrict the concept of political feasibility to just one function in the future (call this the "narrow" approach). So, for example, we might decide to retain the idea of feasibility when we want to allocate blame, but use other terms when we want to make prudential or strategic decisions (rather than calling these "predictive feasibility judgments") or when we want to decide on deliberation-worthiness. On this view, embracing a multifunctional account of feasibility will not add clarity to our deliberations but is likely to make things less clear overall.

I doubt that a narrow approach to feasibility will be clearer, for two reasons. First, if we were to adopt a narrow approach with the aim of achieving greater deliberative or conceptual clarity, the best way to do this involves being clear about what we are narrowing from. In other words, if the label of a feasibility judgment is to be applied only to a specific function in a way that restricts how we have tended to use the word up to now, isolating this function requires us to clearly distinguish it from those functions that have previously laid claim to being feasibility judgments given that these other functions (regardless of what they are now to be called) will continue to play important roles in political deliberative practices. In order for us to re-label these important functions which are no longer to be regarded as "feasibility" judgments, we would need to identify the fact that they currently lay claim to that label, as I do in this article. So, even if one is convinced that the clearest

[27]I am grateful to an anonymous reviewer for pressing me on this point.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

approach to feasibility is a linguistically narrow approach, a clear conception of feasibility will still need to recognize and take into account the kinds of functions I identify here (and perhaps others as well), even if one decides to attach different labels to them while retaining "feasibility judgment" for only one of these functions.

If I am right that a linguistically narrow approach still needs to do the work of identifying the various functions I discuss in this article if we are to to properly identify an appropriate narrow conception of feasibility, it is not clear that this process is any simpler or more straightforward than an expansive one (such as my preferred account), given that each approach still ends up with a plurality of judgment types.

Supposing that I am wrong about this, a further practical problem is that there is no obvious candidate function with a better case than the others to retain the label of "feasibility judgment"—the various functions in question all seem to play important roles in public political deliberation (which is what we would expect given that monofunctional accounts try to capture intuitive understandings of what feasibility means). Of course, it is important not to beg the question here given that we could interpret the work of various theorists discussed in the paper as attempts to convince us that their preferred function is the right one, but the existence of these debates shows at the very least that there is probably some reasonable disagreement to be had over the proper use of the term which suggests that it will be at least a practical challenge to persuade people to change how they talk about feasibility. This is a challenge for a narrow approach that is not faced by an expansive one, which does not require us to substantively change our ordinary practices in talking about feasibility, but only that we be more precise about what we are trying to do with feasibility judgments in particular cases. This kind of clarity is unlikely to be especially onerous given that (as I mentioned above) our ordinary deliberative practices already include some important distinctions (if often implicitly) such as the distinction between scalar and binary feasibility and political and non-political feasibility. Being a little clearer about what kind of feasibility judgment we are deploying seems less demanding than convincing people in general to restrict the term to only one kind of judgment.

Nevertheless, a critic might be willing to accept the move from a monofunctional to a multi-functional approach for the reasons set out above, but insist that this still leaves too much unsaid when it comes to deciding between different candidate conceptions that target the same function. For example, in this article, I have taken both Estlund and Stemplowska's accounts to be examples of feasibility as blame-allocation, but if we are to use feasibility judgments for this function, we still need to decide whether we should deploy Estlund's version, Stemplowska's version, or something else. A critic may suggest that the action-guiding promise of a functionalist approach falters unless we are able to render more specific judgments about which conception to adopt for which function.[28]

The account I have developed here suggests at least two preliminary steps we should take when settling on the right conception for some particular context. First, we need to be clear about the relevant function for the context in question—are we interested in allocating blame, making a prediction, determining deliberation-worthiness, or something else? This initial step will at least help us identify the right "family" of conceptions from which to choose (for example, if we are interested in allocating blame, we can identify Estlund and Stemplowska's views as relevant candidates, while excluding conceptions like Southwood's or Wiens'). The next step is to consider the extent to which each conception under consideration succeeds in performing the relevant function. In the case of blame-allocation, this seems to be a matter of working out which conception does the best job of capturing our intuitions about blame. If we were considering feasibility judgments as predictions, however, the test would be quite different, since the best predictive account is the one that most accurately predicts the success or failure of political proposals. I will not

[28]I thank an anonymous reviewer for raising this point.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

oversell the multifunctional approach here—it is not a panacea since we still need to do the hard work of deciding between specific conceptions that aim to serve the same function—but such an account can nevertheless provide a degree of clarity that will help frame the initial investigation, and makes it less likely that we will dismiss good concepts for bad reasons (for example, by misapplying accounts that aim to serve one function to tasks for which they are not appropriate). Ultimately, the only significant challenge for this kind of pluralism is that we must be careful to be clear about what conception of political feasibility we are using and why, but this seems a small price to pay in exchange for the flexibility a multifunctionalist approach provides.

6. Conclusion

The central contentions of this paper are that we ought to reject monofunctional accounts of political feasibility in favour of a multifunctional account and that we should recognize and utilize feasibility judgments for at least those functions discussed above (i.e. to determine deliberation-worthiness, to allocate blame, to make predictions and to make comparisons). It must be emphasized that this is not intended to be an exhaustive list—I leave open the possibility that other feasibility functions may be identified or that there may be more helpful ways to conceptualize the functions as described above, but I remain sceptical that any functionalist account will be able to condense all of our deployment practises into a single function.

Two important implications seem to follow from this approach. First, if we accept that a multifunctional account is best, this splits subsequent theorising into two distinct but related tasks: (1) identifying all of the relevant functions that feasibility can serve (such as those discussed in this article), and (2) making sure that the appropriate kind of feasibility judgment is being deployed to achieve one's aims in some particular context.

The second important implication is for political activists, who can use the categories developed here as part of a framework for working through the ethics of political strategy. For instance, where activists understand themselves to be deploying predictive feasibility judgments this categorization can help to reveal which kinds of empirical facts should have a bearing on their decisions (that is those affecting actual probabilities) and which ethical principles should constrain them (for example, those concerning the ethics of risk). Similarly, where activists' projects do not depend on the accuracy of their predictions (for example, because they are deploying feasibility judgments to allocate blame or determine deliberation-worthiness), these categories provide a way to talk about feasibility without being vulnerable to charges of utopianism or naiveté. Ultimately, however we carve up the conceptual terrain and whether we take the perspective of theorists or activists, a multifunctional account of political feasibility seems to provide greater theoretical and practical flexibility at a negligible cost.

Brian Carey is an Assistant Professor in Political Theory at the School of Government and International Affairs, Durham University. His research focuses on nonideal theory in contemporary moral and political philosophy, especially issues concerning public deliberation, linguistic justice, and political feasibility.

References

Brennan, G., & Southwood, N. (2007). Feasibility in action and attitude. In J. Josefsson, & D. Egonsson (Eds.), *Hommage a` Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. Department of Philosophy, Lund University.

Carey, B. (2015). Towards a 'Non-Ideal' non-ideal theory. *Journal of Applied Philosophy*, 32(2), 147–162.

Carey, B. (2016). Justice for Jerks: Human nature, selfishness, and noncompliance. *Social Theory and Practice*, 42(4), 748–766.

Doroschuk, I. (1982). *The safety dance. Rhythm of youth*. Virgin Records.

Ermann, E., & Moller, N. (2020). A world of possibilities: The place of feasibility in political theory. *Res Publica*, 26, 1–23.

Estlund, D.. (2008). *Democratic authority: A philosophical framework*. Princeton University Press

Estlund, D.. (2011). Human nature and the limits (if any) of political philosophy. *Philosophy & Public Affairs*, 39(3), 207–237.

Gilabert, P., & Lawford-Smith, H. (2012). Political feasibility: A conceptual exploration. *Political Studies*, 60(4), 809–825.

Southwood, N. (2016). Does "Ought" imply "Feasible"?. *Philosophy & Public Affairs*, 44(1), 7–45.

https://doi.org/10.1017/can.2025.10013 Published online by Cambridge University Press

Southwood, N. (2022), "Feasibility as Deliberation-Worthiness", *Philosophy & Public Affairs*, 50(1), pp. 121–162.
Stemplowska, Z. (2008). What's ideal about ideal theory?. *Social Theory and Practice*, 34(3), 319–340.
Stemplowska, Z. (2016). Feasibility: Individual and collective. *Social Philosophy and Policy*, 33, 273–291.
Stemplowska, Z. (2021). The incentives account of feasibility. *Philosophical Studies*, 178, 2385–2401.
Valentini, L. (2012). Ideal vs. non-ideal theory: A conceptual map. *Philosophy Compass*, 7(9), 654–664.
Wiens, D. (2015). Political ideals and the feasibility frontier. *Economics and Philosophy*, 31(3), 447–477.

Cite this article: Carey, B. 2025. A Multifunctional Account of Political Feasibility. *Canadian Journal of Philosophy*: 1–15, doi:10.1017/can.2025.10013