Understanding Political Feasibility[1]

Holly Lawford-Smith

[Last edited 18[th] June, 2012. Forthcoming in *Journal of Political Philosophy*].

I. Introduction

We only have to go back fifty years to find women conditioned into the domestic servitude of men; and a few hundred years to find black people kept by white people as slaves. At either of these historical junctures, someone might have asked: 'is emancipation feasible?' 'Is gender equality feasible?' We know that the answer to both questions is 'yes', because we have borne witness to the relevant social changes. Slavery has been abolished, and equal rights for women have been won. Inequalities deriving from race and gender persist, but in nowhere near as severe a form.

In the present, we find industry engaged in factory farming on a massive scale, resulting in the torture of a great number of animals; we find many nations with nuclear weapons, resulting in the constant possibility of nuclear warfare; we find sea levels rising rapidly due to climate change, with the result that low-lying nations are at threat of being completely submerged. We might ask: 'Is abolishing factory farming feasible?' 'Is universal nuclear disarmament feasible?' 'Is a radical reduction in the global concentration of greenhouse gases in the atmosphere feasible?' We do not know the answers to these questions, but we can make more or less educated guesses.

This is a paper about those more or less educated guesses. Political philosophers often criticize one another's theories on the grounds that they are infeasible. As Juha Räikkä has pointed out, 'the notion of political feasibility has always been politically significant. There has always been some tendency to reject new and unbiased suggestions on the grounds that they are "impossible" or "idealistic" or "utopian". This is why it is practically important to be aware how the notion of feasibility is used and should be used in political arguments that represent

political theory'.[2] I'm interested in what the criticism of infeasibility actually amounts to.[3]

The account of feasibility that I'll present in this paper has two main implications, one theoretical and one more practical. On the former, those who have written explicitly about feasibility have preferred a binary view roughly equivalent to the 'ought implies can' constraint in moral theory.[4] But that view makes feasibility all but useless in politics. We don't want to rule options off the table entirely – rather, we want a way of saying how feasible some alternative is, which is just one consideration that feeds into our deliberation about what, ultimately, we ought politically to do. So one of my aims in this paper is to distinguish binary from scalar feasibility and defend a much greater focus on scalar feasibility within political theory, because it is scalar feasibility which has an active role to play in real political decision-making, and scalar feasibility which determines the extent to which a political theory is ideal or non-ideal (which collapses the standard assumption that a theory is one or the other and somehow problematic in virtue of being ideal). On the latter, those who actually make or are interested in making feasibility assessments have tended to be insufficiently sensitive to the importance of making such assessments in an agent-relative way, and the diachronic aspect of feasibility. On my account, both of these features are prominent, which results in an account that I think will actually be more helpful for making feasibility assessments in real politics.

In §II I take a closer look at the role that feasibility plays within political theory. In §III I ask about the structure of claims about feasibility, as a way of drawing out its important features, and in §IV I distinguish binary from scalar feasibility and defend, against those writing about feasibility thus far, the scalar view as more interesting and more useful, both in political theory and in political practice. §V concludes.

---

2 Räikkä 1998, p. 39.

3 Räikkä (1998) distinguishes 'political feasibility' from 'the feasibility condition in political theory', arguing that a theory can meet the feasibility condition without being *politically* feasible, by which he means that the political will to bring the theory's recommendations into practice might be lacking. In his terms, I am concerned in this paper with 'the feasibility condition', and I try to talk in terms of 'feasibility constraints' or 'requirements of feasibility'. Any usage of 'political feasibility' should however be read as synonymous with those, rather than as indicating an interest in distinctively political constraints.

4 Gilabert & Lawford-Smith (2011) acknowledge the scalar alternative, but do not defend the latter as more interesting or important than the former.

II. What role does feasibility play in political theory?

A strong feasibility requirement is that every political theory[5] must take seriously certain facts about how the world is. For example, a theory that requires all citizens to be provided with more than the average national income is infeasible because it violates a constraint that every theory must respect, namely the constraint of logical consistency.[6] If there really is such a requirement, then the project for those interested in feasibility is to figure out just what those facts are which every political theory must take seriously.

A weaker feasibility requirement is that some political theories must take seriously some facts about how the world is. A good candidate for the kind of theory bound by such a requirement is theory issuing more or less immediate imperatives to people. For example, a theory about animal liberation should issue the imperative that people stop eating factory-farmed meat only if the world is such that people can stop eating factory-farmed meat; a theory about nuclear disarmament should issue the imperative that nations multilaterally disarm only if nations can multilaterally disarm (of course, the challenge is to say what 'can' consists in here).

I think both the strong and the weak feasibility requirement have a place in political theory. The strong requirement is true in that some extremely general constraints bear on all political theories, no matter where they fall in the spectrum from ideal to non-ideal.[7] The weak requirement is true in that the more non-ideal a theory is, the more stringent are the constraints that bear upon it. In this paper, however, I want to concentrate not on the different constraints that apply to different kinds of theories,[8] but on two different ways of thinking about feasibility.

Feasibility can play several roles in political theory. Usually, it is employed as a tool for ruling out theories that cannot be implemented. But I will argue that it is more usefully employed as a tool for ranking alternative theories along one of the dimensions relevant to making decisions about what to actually do. Those are the

5 I use 'political theories' as a placeholder, pending discussion in the next section about the objects of claims about feasibility (candidates include political theories, the requirements issued by political theories, actions, sets of actions, outcomes, states of affairs).
6 I owe this example to Pablo Gilabert.
7 See discussion in Miller 2008, and Pogge 2008. Elsewhere I have argued that ideal theory has much more latitude in failing to take feasibility constraints seriously than non-ideal theory does, and that ideal theory is not 'bad' theory for that reason. See Lawford-Smith 2010.
8 On that topic see discussion in Gilabert & Lawford-Smith 2011, esp. 10-15.

binary and the scalar roles, respectively. It can be used to reveal the powers that certain things have (e.g. to undertake certain actions or bring about certain outcomes); if φ is feasible for an agent, then the agent has the power to φ. It can be used to supplement practical reasoning about action choice. We decide what to do by weighing the probabilities[9] of outcomes against both their desirability and the risks inherent in the actions intended to produce them. Feasibility is independent of both desirability and risk, but understanding it is part of understanding the conditional probabilities that are part of the decision-making process. Related to this last role, feasibility can be used as a heuristic, in combination with rough considerations about value, when accessing the full decision-theoretic calculus is impractical. Decision theory outputs the optimal action only after a protracted set of calculations, which we often won't have the time or the energy for (although we'll admittedly have more time and resources for making elaborate calculations when it comes to deciding about the actions of the state than when making everyday decisions as individuals). We have an intuitive or antecedent sense of what is desirable – often we will have settled on a desirable outcome in advance, and just be wondering which of the available actions to pursue in order to bring it about, or to bring it about most efficiently. We can trade this off against an idea of what is feasible. Understanding feasibility better will allow us to do this.[10]

These are at least five important roles for feasibility to play in political theory; yet feasibility has so far proved to be fairly elusive and under-theorized.[11]

III. The structure of claims about feasibility

In this section I ask about the structure of claims about feasibility, in order to shed light on the elements that are crucial to understanding them. What kinds of things

9 As I will go on to discuss, these are conditional probabilities – the probabilities of outcomes given the relevant actions.

10 As a heuristic, this is not perfect; leaving risks out of the calculation means that sometimes the heuristic will give the wrong answers. That's okay so long as the convenience of the heuristic overall outweighs the cost of the errors. I am grateful to Geoffrey Brennan for discussion on this point.

11 For the discussion on feasibility see: Barry & Valentini 2008; Brennan & Pettit 2005; Brennan & Southwood 2007; Cowen 2007; Gilabert 2009; Gilabert *forthcoming*; and Raikka 1998. For the discussion on what is to all intents and purposes the same issue see: Buchanan 2004; Cohen 2009; Estlund 2008, ch. 14; Estlund 2011; Hawthorn 1991; Wiens, ms.; and Jensen 2009. I do not include the growing literature on ideal and non-ideal theory here because while it is relevant to the related project of what kinds of theories feasibility constraints apply to, it is not related to the main concern of this paper, namely, understanding feasibility itself.

are feasible and infeasible? Whose actions matter for feasibility assessments? What does it take for the relevant object to be feasible? Does context matter? Is feasibility indexed to particular agents, or not? To keep things relatively simple, I'm going to talk about feasibility in this section as though it is binary (something is feasible, or not). Then in the next section, I'll explain how to think about the account in a way that is scalar instead.

Claims involving feasibility typically go like this: 'it's feasible that the global concentration of greenhouse gases in the atmosphere be maintained at levels necessary to prevent the sea-level rise projected to submerge low-lying countries'. They are claims about outcomes, or states of affairs. Different normative / political theories issue different imperatives, for example global egalitarianism applied to the *status quo* recommends substantial global redistribution of wealth. These imperatives exert normative pressure on agents to bring about the relevant outcomes. If we think that such theories are bound by feasibility constraints, then we'll presume that they're claiming that a certain outcome, desirable by the theory's own lights, is feasible.[12] The challenge is to say what makes such a claim true or false; to say what feasibility consists in. Let's start with the following commonsense claim:

(a) An outcome is feasible *iff* it can be brought about.

That much is fairly obvious. Brought about by what, or whom? Outcomes are generally the result of the actions of agents. Of course there are exceptions: the outcome in which my house has fallen into disrepair and my garden has become choked by weeds is not the result of any action of mine – rather it is the result of an inaction, or omission; the outcome in which the sun has risen tomorrow is dependent on neither the action nor inaction of any agent. We can account for the former kinds of cases by linking outcomes to agents, whether by action or inaction, and we can bracket the latter kinds of cases as largely irrelevant to feasibility assessments. It's unlikely that a normative theory will issue the imperative that the sun rise, so we're unlikely to need to test the outcome in which the sun rises for feasibility. Which takes us to:

---

12 To keep the discussion simple I'll just talk in terms of outcomes, but I am concerned in particular with the outcomes that the imperatives of normative theories direct the agents on whom they bear towards. It actually looks like whole normative theories can be feasible or infeasible too: if a theory issues a finite number of imperatives, and we know all of them to be infeasible, then we can reject the theory as infeasible too (although see later discussion about diachronic feasibility).

(b) An outcome is feasible *iff* there exists an agent who can bring it about.

It is immediately apparent that 'agent' cannot be limited to individuals. The feasibility of the vineyard's grapes being picked before they go bad cannot depend on the superhuman speed and dexterity of one individual. Surely such an outcome is feasible if there is a set of individuals who could pick the grapes together, or a collective agent (such as a fruit-pickers' union) which could see to the grapes being picked. Be we have to be careful when we talk about what a set of individuals can do together. It's true that a set of individuals might agree to pick the grapes together, but it's equally true that there might be one agent willing to pick grapes with others, and the remaining agents completely unwilling. In the former case we'll want to say it's feasible that the grapes be picked, in the latter not. What differentiates the two cases is the extent to which the individuals are 'collectivized' in the sense that they have a group-level decision-making procedure which allows them to coordinate the actions of their members.[13] Thus 'agent' should be taken to include both individuals and collectives, and exclude uncoordinated aggregates of individuals.

Still, we need to know more about what it takes for it to be true that an individual or collective agent can bring about an outcome. One way to fill this in is to say that the relevant agent must have as one of her (its) options an action with a positive probability[14] of bringing the outcome about.[15] That just leaves the slightly more tractable problem of establishing when an agent has an action in her option set and when not. Thus:

(c) An outcome is feasible *iff* there exists an agent with an action in her (its) option set[16] that has a positive probability of bringing it about.

---

13 See discussion in Collins ms., p. 9, and references therein.

14 By 'positive probability' I mean non-zero probability. I'm setting aside problems with infinities here (when an action has infinite possible outcomes we have to assign them probability zero, because if we assign them any positive probability then they won't together sum to one) on the assumption that such problems won't generally plague political feasibility assessments.

15 In general, options are actions, i.e. the set of options an agent has is just the set of actions available to her. But action theorists allow that 'options' can also include the outcomes produced by the actions available to the agent. So for example, one of my options (actions) is to write this footnote, and another of my options (outcomes) is to have written this footnote. So to avoid ambiguity on this point I'll talk in terms of actions in an agent's option set rather than options in general.

16 For a logic of feasibility to be possible, we must include options of the form 'be such that...'. That is to allow statements like 'it is feasible for John to be such that $2 + 2 = 4$', or 'it is feasible for Harry to be such that it is raining'. These are true so they are actual, and if they are actual then *a fortiori* they are feasible (at least in the binary sense) but they are in the same class as the action- (or omission-) independent outcomes bracketed above.

This seems to get things roughly right, supposing we can say which actions are in agents' option sets. To illustrate, imagine we're interested in a normative theory which issues the imperative 'all the grapes in the vineyard must be picked'. If that theory is true, the imperative bears on individuals and directs them toward an outcome, one in which all the grapes in the vineyard have been picked. We want to know whether this outcome is feasible, because we believe that the imperatives of normative theories (and in some cases the theories themselves) are subject to feasibility constraints. So we ask whether there is some agent, individual or collective, who (which) has an action in her (its) option set that has a positive probability of bringing it about that all the grapes in the vineyard are picked.

Suppose that there is in fact a fruit-pickers' union working in the vicinity of the vineyard. It is not otherwise engaged, so we take it to have an action in its option set of picking the grapes. The action of picking the grapes has a positive probability (in fact a very high one) of resulting in the outcome that the grapes are picked. Thus the outcome the theory directs agents towards is feasible. Suppose on the other hand that there's only one individual in the vicinity of the vineyard. She is not otherwise engaged, and she surely has the action in her option set of picking at least some of the grapes. But in this case, no matter how hard she works, she simply won't be able to pick all the grapes before they go bad. If we're liberal about credence, then we might want to assign her action probability zero, in which case we can say that the outcome in which all the grapes are picked is infeasible, because there's no agent with an action in her (its) option set with a positive probability of bringing it about. If we're conservative about credence (preferring to assign a probability of 0 or 1 only when absolutely certain) then we might want to assign her action some low probability, in which case we can still say that the outcome is feasible (in the next section I'll show how treating feasibility as scalar rather than binary gets around the current inability of the account – assuming conservativism about credence – to distinguish the single-individual case from the fruit-pickers' union case in terms of the feasibility of the outcome).

There's still a problem to account for, and that is that we can't say decisively that an outcome is infeasible if it isn't fairly explicit what the time and place of its being brought about is meant to be. Thing were simple in the vineyard case, because

we have some idea of how long it takes for grapes to go bad. But what about very general imperatives such as 'wealth ought to be redistributed from the richest countries to the poorest'. If the three richest countries have the action in their respective option sets of redistributing wealth to the three poorest countries, but the three next-richest countries do not, should we say that the outcome in which global wealth has been redistributed from richest to poorest has been realized, or not? What if none of the richest countries right now have that action in their option sets, but they'll have it in twenty years? In that case should we say the outcome is feasible, or not?

The geographical question is easier addressed than the temporal question. Consider the case in which there are no individuals or collectives living in the vicinity of the vineyard mentioned above. What if the nearest individuals all lack an option of picking the grapes alone, and the nearest collective with the right kind of action in its option set is several countries away. Should we say the outcome is feasible just because there's an agent somewhere in the world with an action in its option set of picking the grapes, even if undertaking that action would mean flying across several countries, at serious cost in terms of time and money? Strictly speaking, the answer is yes. Sometimes we do fly experts across the world to undertake actions that no one else can perform (think of medical specialists, neurosurgeons, forensic scientists, and so on). An outcome can be feasible and yet not worth trying to bring about, because of the cost involved. Feasibility assessments are empirical assessments. We want to know whether an outcome can be brought about. If there's an agent somewhere in the world who has an action in her (its) option set with a positive probability of bringing the relevant outcome about, that outcome is feasible.

The temporal question is more difficult, because it involves diachronic possibility rather than merely synchronic possibility. What agents can do later depends in part on what they do now. I take this distinction between synchronic and diachronic possibility from the discussion of abilities, which distinguishes abilities agents have at a given time, and abilities they can come to have at a later time depending on which of the abilities they have at the given time they choose to exercise. For example, I cannot speak Portuguese, but I have the synchronic ability to enroll in a Portuguese language class, and thus I have a diachronic ability to speak

Portuguese – I am able to speak it at some later time.[17] It is plausible that when we ask about what's feasible, we'll want to include at least some actions that are diachronically but not synchronically in agents' option sets. For example, we might ask about the feasibility of a certain set of Australian citizens voting in the national elections. Imagine that the citizens are not registered to vote, but that it's easy to register (one simply fills out an online form). It would be strange if we had to say that their voting was infeasible, because for each of them, they did not have an action in their option set at the time that gave their having voted a positive probability. A partial solution to this problem is to include a temporal index:

> (d) An outcome is feasible *iff* there exists an agent with an action in her (its) option set within the relevant temporal period that has a positive probability of bringing it about.

This makes things easier in cases where we are fairly sure about what will be the case at some later time. For example, imagine that registration to vote is facilitated by giving the relevant citizens' a paid half-day off work and making sure they all have access to the internet, and imagine that the relevant citizens are actually interested in voting and want to make use of that facilitation. Then we can be fairly sure that they will register to vote, so we can be fairly sure that they will, within the specified temporal period, come to have the synchronic ability to vote. In this case we are not concerned so much with diachronic abilities as with synchronic abilities over time – we're interested not in what the agents could come to be able to do, but what they will be able to do. That is why the solution is only partial.

The more complicated element is how to think about the actions that agents merely could come to have in their option sets, where there's no likelihood that they will come to have them within the relevant temporal period, and in particular actions that agents could come to have not merely in virtue of their own choices but in virtue of the choices of others. One action that is sometimes available to individuals is to collectivize – but the availability of that action is contingent upon the willingness of at least one other individual. If agents were to undertake that action, assuming the willingness of at least one other agent, then there would be a

17 See discussion in (A. Goldman 1970, pp. 204-204), (H. Goldman 1976, p. 453), (H. Goldman 1978, p. 193 & see also references in fn. 17), (Jensen, 2009), (Gilabert, 2009), and (Gilabert & Lawford-Smith, 2011).

collective agent, which in turn would have an option set with all sorts of actions in it. And that collective, like any extant collective, would have (has) as one of its options the extension of its powers, the addition of new members, or the combining together with other collectives, all of which would increase the number of actions in its option set.

What we want to avoid is having such a permissive account of available actions that outcomes like ending global poverty or achieving global carbon neutrality come out as feasible. In those cases the relevant agents are countries. It is true of every country that there are actions they could take now that would lead to their having an action available to them later on that would, in combination with other countries doing the same thing, lead to carbon neutrality or global sufficiency. Ending global poverty and achieving global carbon neutrality are both *possible*. But we don't want to say these things are feasible, because we don't want to count all merely possible actions as available in the relevant sense, and we certainly don't want to ignore the importance of collective action problems for infeasibility. The full solution to the problem, I think, comes from insisting on agent-relativity in feasibility assessments.[18] By that I mean that it is very important that we distinguish the agent whose option set we're interested in from all the other agents upon whose actions the outcome might depend. I develop this idea in more detail in the next section, in discussing scalar feasibility.

I have said that an outcome is feasible if there is some agent whose action has a positive probability of bringing it about. But consider the difference between 'it is feasible that Jonathan run 10 kilometres' and 'it is feasible for Jonathan to run 10 kilometres'. What it takes for each to be true is quite different. According to (d), the former is true so long as there is some agent who has an action available to him/her/it that could bring about the outcome that Jonathan runs 10 kilometres. Suppose that Stephanie has the action available to her of threatening to spike all his

---

18 As part of the related discussion about 'ought implies can', Frank Jackson suggests that only oughts that bear upon individuals (of the form '*X* ought to *A* at *t'*) entail ability; oughts that bear on states of affairs (of the form 'it ought to be that *X* does *A* at *t*) do not. See (Jackson, 1985). I think that *some* oughts that bear on states of affairs are subject to feasibility requirements (the ability of some agent to bring them about), namely the non-ideal ones. But perhaps this is because of the connection to agent-relative oughts, namely that part of the success conditions of a feasible-that claim are that at least one feasible-for claim is true (see discussion below).

vegan meals with dairy products, which would be effective in coercing him into running the 10 kilometres. In such a case, Jonathan running 10 kilometres is feasible because Stephanie has an action available to her that could bring it about. But that is certainly a bit strange. Shouldn't we rather care about what actions are available to Jonathan, and whether one of them has a positive probability of bringing about the outcome that he run 10 kilometres? That will depend on the normative imperative we're assessing for feasibility. If it merely directs agents toward an outcome in which Jonathan runs 10 kilometres, then it seems perfectly right that we should be interested not only in what actions are available to Jonathan himself, but also what actions are available to other agents – what others can do to get Jonathan to run the 10 kilometres. But the imperative might be issued by a theory in which it is important that Jonathan run the 10 kilometres for his own reasons and not because he is coerced by others (as it will be in any theory which cares not just about consequences but also about virtue / motivations / intentions).

IV. Distinguishing binary and scalar feasibility

We can formalize (d) in the following way:

$$BF(O) \text{ iff } \exists A \, P(O \mid A) > 0$$

Which says, an outcome (O) is binary-feasible (BF) if and only if there exists an action ($\exists A$) such that the probability of the outcome given that action is greater than zero. Actions belong to agents, so there only exists an action when it is had by an agent, in her option set. So 'there exists an action...' is shorthand for 'there exists at least one agent with an option set containing an action...'. Two crucial and related questions remain. The first is: when does an agent have some action as an option? If we cannot establish when an agent has an action in her option set and when she does not, then then the definition of feasibility in (d) will be useless. The second is: how should we think about others' actions, in relation to what is in a given agent's option set? Normally an action is something that is under the full control of an agent. But sometimes the actions of others are partially under my control, or depend on my action. I'll address each of these questions in turn.

Others have suggested that there are two dimensions involved in feasibility:[19]

---

19 In the extant literature, all writers assume the binary view. See (Gilabert & Lawford-Smith, 2011)

'accessibility', and 'stability'.[20] An action is accessible to an agent if she could choose to do it; an outcome is accessible to an agent if one (or a set) of her actions could produce it. Notice I say only that the agent *could* choose the action, and that the action *could* produce the outcome. This is a fairly weak understanding of feasibility, but one that is suitable for the binary role in which feasibility is used to decisively rule out recommendations that cannot be implemented. An action is not decisively ruled out so long as the agent could in principle choose to do the action (even if she almost certainly will not), and so long as that action could in principle realize the outcome (even if it almost certainly will not).

An action is only a means to an outcome if it would produce that outcome in a way that would be more or less stable.[21] Geoffrey Brennan and Philip Pettit stress the importance of stability considerations, arguing that political philosophers should be more concerned with 'the problem of how to ensure that whatever arrangements are put in place ... are arrangements that ordinary humans are able to sustain'.[22] Gerald Cohen also emphasizes the importance of stability, and suggests that the crucial question is not whether people can change, but whether we can change institutions in a way that they can change people.[23]

Even with the concepts of accessibility and stability in hand, we still need to know what makes it the case that an action is accessible to an agent, or that a stable outcome is accessible by way of an action. What kinds of facts settle these questions?

The kinds of facts that settle binary feasibility can be referred to as 'hard constraints', as opposed to the 'soft constraints' that bear upon scalar feasibility, which I'll come to soon. Hard constraints include facts about what is logically, conceptually, metaphysically and nomologically impossible, and these serve to limit the option sets available to agents. No agent can do the nomologically impossible.

---

for a recent survey.

20 Cohen 2009, Cohen 2001, and for a different version of accessibility see Buchanan 2004.

21 Some might prefer to leave stability out of a conceptual account of feasibility, and in that case it could be repackaged as part of the desirability of some action. It is undesirable to put our efforts into pursuing an outcome that won't last (unless it is a transitional outcome). But it seems to me that the recommendations of political theories come with an implicit or explicit temporal scope, so that when we ask whether an outcome is feasible, what we're really asking is whether an outcome, for a given length of time, is feasible. The answer will be 'yes' if the outcome, for that given length of time, meets the conditions for binary feasibility.

22 Brennan & Pettit 2005, p. 264.

23 Cohen 2009, pp. 56-57, Cohen 2001.

Hard constraints rule actions out of option sets, and when actions are not in option sets the outcomes that rely on them are infeasible in the binary sense. The kinds of impossibilities just mentioned are timeless in that a theory cannot violate them and be feasible. But there are other kinds of hard constraints, things which are impossible now but may not be impossible in the future. For example, the absolute limits of a budget, or the absolute limits of the available workforce, are a hard constraint upon a construction project (by 'absolute' I mean to include the usual range of beg- steal- and borrow- ings). These are time-sensitive hard constraints. There are agent-relative hard constraints of both kinds. For example, it is a timeless hard constraint upon an incurable paraplegic that she cannot do things that require mobility in the lower limbs, and it is a time-sensitive hard constraint upon the monolingual that they cannot speak in a language other than their mother tongue.

A political theory is hardly likely to make recommendations that violate timeless, non- agent-relative hard constraints, as others have noticed.[24] But a theory may well violate time-sensitive hard constraints, because of a lack of information about resource availability. There are people who think that we ought to introduce artificial rain clouds over the Sahara Desert as a way of making the land arable again.[25] Whether or not this is feasible depends on the technologies capable of generating artificial rain clouds, and the resources required to implement such a large-scale project. And a theory may also violate agent-relative hard constraints, because of a lack of information about the agent and her current capacities.

We can incorporate the notion of hard constraints (including both timeless and time-sensitive, agent-relative and non- agent-relative) to give the truth conditions for the definition of feasibility given in the last section. The definition was:

> (d) An outcome is feasible *iff* there exists an agent with an action in her (its) option set within the relevant temporal period that has a positive probability of bringing it about.

And we can give the truth conditions as the following:

> An agent has an action in her (its) option set *iff* her performing that action is not

24 See e.g. Jensen 2009.
25 On cloud seeding see http://www.newscientist.com/article/mg18624952.000-kicking-up-a-storm-with-the-cloud-seeders.html; on the laser version of the technology see http://www.newscientist.com/article/dn18848-laser-creates-clouds-over-germany.html.

ruled out by any hard constraint. An action has a positive probability of (stably) bringing about an outcome *iff* the outcome being produced (stably) by that action is not ruled out by any hard constraint.

These truth conditions incorporate conditional probability. An agent has an action in her option set when that action has a positive probability of being performed given the agent's choice (intention / trying / willing) to perform it; and an outcome is accessible by way of an action when the outcome has a positive probability of becoming actual given the action.

As mentioned earlier, feasibility in this binary sense has a lot in common with the principle that 'ought implies can', much-discussed in moral philosophy. The idea is that the proposition 'X is obliged to φ' is false if X cannot φ, on some suitable understanding of 'cannot' such as 'lacks either the ability or the opportunity to'.[26] By contraposition, 'ought implies can' gives 'not-can implies not-ought', or, more colloquially, 'if it's not the case that a person can φ then it's not the case that she ought to φ'. But we do not want these constraints to be too strong, or they'll rule out oughts that shouldn't be ruled out. This is my main reason for thinking we should be more interested in scalar feasibility than binary feasibility. If there is a requirement that normative imperatives (or whole theories) are feasible, then whichever imperatives do not meet that requirement will be ruled out (whether as imperatives / theories, or merely as bearing on agents, is not clear). We don't want to rule out the possibility of realizing really good outcomes unless we're fairly sure those outcomes are unrealizable, so we should make the binary feasibility constraint weak. But then it is unlikely to do much work.

Surely the better thing to say is not that feasibility is required, but that feasibility matters. It matters but so do other things, like the desirability of an outcome, or the risks inherent in pursuing it. Scalar feasibility allows us to say how feasible outcomes are, and then we can use that datum against the other relevant considerations in deciding what to do. Sometimes it will be worth pursuing an outcome with low scalar feasibility, because having brought it about would be really good, and sometimes it won't be worth pursuing an outcome unless it has high scalar feasibility, because having brought it about won't make all that much difference to

---

26 For a recent survey see Vranas 2007.

the goodness of the world. It's never worth pursuing an outcome with zero feasibility, which is how binary feasibility still has a role to play – just not a central role.

In the scalar role, the recommendations of theories can be ranked according to how feasible they are. The kinds of facts appealed to to determine this ranking must be facts that make an outcome conditionally less likely to obtain, rather than facts establishing that it cannot obtain. These are the 'soft constraints' mentioned earlier.

There were two parts to the truth conditions for binary feasibility: a positive probability of an agent undertaking a particular action if she chose to, and a positive probability of an action's producing a particular outcome. For the scalar definition there's only one part. We use a conditional probability in that we *assume* the agent's choosing a particular action.[27] Importantly, we do not assume *all* agents' choosing of particular actions, and we do not assume non-agents' (e.g. non-collectivized aggregates of persons) choosing of particular actions. But we do assume the trying of the agent we're interested in. Why? Primarily because we don't want to let agents off the moral hook. The fact that a person won't do something isn't enough for us to retract an imperative that she ought to. It's the same for actions an agent is unlikely to do. A person's preferences over her own actions do not determine the limits of her obligations. What matters is the extent to which the action in her option set is likely to produce the relevant outcome. The more likely the outcome given the action, the more feasible the imperative issued by the theory; the less likely, the less feasible. We can formalize that in the following way:

$$SF(O) = P\,(O \mid \max_A)$$

Which says, the scalar feasibility (SF) of an outcome (O) is equal to the probability of the outcome given the best (or best equal) action. Again, actions belong to agents, so 'the best (or best equal) action' is shorthand for 'the best (or best equal) action in the option set of any agent'. Talking in terms of the best action is a way of being charitable about scalar feasibility. We don't ask about just any action; rather we ask about the action that gives the outcome the best chance of coming about. It is the difference between considering the chance of a window breaking given my throwing a brick through it, compared with the chance of it breaking given my merely staring

---

27 In this I follow Brennen & Southwood 2007.

at it. (bracketing special arrangements where my stare is the secret signal for another person to throw a brick). What kinds of facts fall under the umbrella of the 'soft constraints' used to determine the probability of the outcome given the best-placed action?

The three most obvious kinds of soft constraints upon an action's bringing about some outcome are economic, institutional, and cultural.[28] Facts about the current economic system make outcomes featuring a different system unlikely to succeed, facts about entrenched political institutions make outcomes clashing with those institutions unlikely to succeed, and facts about religion and culture make outcomes featuring different beliefs and attitudes unlikely to succeed (let culture extend also to the constraints posed by the positive morality of a society). Conflicting with any of them does not make an outcome infeasible in the binary sense, but it might make it less feasible in the scalar sense. For example, it is not impossible to raise support for socialist reforms within a capitalist economy, but the fact that the economy is capitalistic, and the extent to which people support the fact that it is so, will make socialist reforms less likely to succeed. Likewise, it is not impossible to introduce reforms that go against the culture or religion of the citizens in a society, but we can expect much less compliance, and much more resistance, when that is the case. In some instances, if we want the reforms badly enough, we will have to be prepared to really manipulate people's incentives in order to secure success.

Accepting soft constraints means accepting that the status quo places some limit upon what we can realistically accomplish. But these are limits we can work around. For example, we might think about how we can introduce changes that will gradually erode the soft constraints, so that at a future time, they will not be constraints at all. (It is one thing to use morally reprehensible forms of coercion to abruptly introduce reforms that clash with deeply-held religious beliefs, it is another to increase education in e.g. the natural sciences in order to gradually draw people out of their religious delirium).

There are less obvious soft constraints, such as the constraints of individuals' psychology and motivation. Do we take 'what people are like', in general, for granted

---

28 In what follows I assume that each of these has equal priority, but it is worth mentioning that they might instead have differing relative strengths.

in thinking about what kinds of recommendations are feasible, or only particular aspects of what people are like? Motivation seems like something we should exclude as a soft constraint. The fact that a person won't do what he ought is no reason to think he cannot do it; the fact that a person is unlikely to do what I want him to do is no reason for me not to try to get him to, at least, if his doing so is very important and I am not choosing between competing actions of similar value. At the extreme end of a continuum of pathologies, things like addiction, compulsion, phobia and illness can make a person unable to act in certain ways. But at the less extreme end of the continuum, these pathologies may be little more than a person failing to try.

We surely do not want to say that the recommendations of one theory are less feasible than another just because people are less likely to try to realize the one than the other. Feasibility is a concept that treads a fine line between possibility, on the one hand, and likelihood, on the other. The feasible does not extend to everything and anything that could possibly be done, because that would leave in too many unrealistic recommendations. But neither does it extend only to what probably will be done, because that would leave out too many aspirational recommendations. That's why feasibility is conditional, in the scalar sense, on choosing (or willing, or effort, or trying, or any other synonym you like). I think the right way to deal with the motivation question, which also helps in dealing with the problem of diachronic possibilities and option sets raised earlier, is to say that the motivation of *other people* is part of the context in which an agent acts, and therefore properly a soft constraint on whether her action will succeed. But *her own* motivation is not something to factor in; when we think about what is feasible for her we think about what she can do, and this depends only on what her options are.

The fact that we get different answers about what is feasible depending on whose projects we are interested in is not contradictory. Usually, we will ask feasibility questions from the perspective of the agent whose project is at issue. Consider the difference between asking what is feasible for Kim Jong-un, and what is feasible for Australia with respect to accepting immigrants from North Korea. It is highly feasible for Kim Jong-un to lift restrictions upon exit, but it is not very feasible for Australia to accommodate North Korean immigrants, because that would require Kim Jong-un to lift restrictions upon North Koreans' exit, and that will not happen

anytime soon.

It can be useful to know what people could do if they chose to, e.g. that Kim Jong-un could lift restrictions upon exit if he chose to. Just as there is a distinction between what is possible and what is feasible, there is also a distinction between what is feasible and what is likely. If we cannot communicate with an agent, have no influence over him, and no reason at all to expect him to act in the ways we find feasible for him, then it would be foolish to rely upon his acting in that way. Knowing that an outcome is feasible because of another's action tells us one thing about likelihood, namely that the chance of the outcome is not zero. But it doesn't tell us more than that. Knowing that an outcome is feasible because of our own action is different – because we can factor that information into a decision to actually undertake the relevant action (and so too with actions of others over whom we have some influence or control, we can choose to exercise that influence or control).

One might worry that a feedback loop is created by the fact that we get different answers about what is feasible depending on whose projects we are asking about. The outcomes open to me via my actions sometimes create obligations for other people. For example, if I want to engage in a collective enterprise, and I strongly believe that you will not do your part, then I will ensure that the outcome can be realized without your contribution. But in that case you are 'off the hook' with respect to an obligation to participate, through your own decision to be the kind of person who cannot be trusted to do so. This might seem problematic to the moral philosopher. But, while you may be blameworthy in having the kind of character that makes you unreliable in collective enterprises, and you may have secondary duties arising from your failure to be drawn into a primary obligation, those are not problems for feasibility. Participating in the collective action will be more feasible for you if it is more likely that you'll succeed if you choose to try; facilitating a successful collective action will be more feasible for me if it does not involve you, because you are unreliable.

Many might find plausible the idea that effort is a soft constraint, i.e. a recommendation is less feasible if it would take more effort to execute.[29] I think that is a mistake. Imagine that an agent has an action in his option set, but that it would

---

29 I am grateful to Clas Weber for discussion on this point.

be really, really hard for him to choose it, or to undertake it once chosen. Is that enough to make the action (and the outcomes it stands to produce) less feasible? No. Brett could put on a performance of his wife's favourite opera in a desperate attempt to save their marriage. This would take a lot of effort, in fact all of his time and resources. But it is an action in his option set, because strictly speaking he has the time and resources. The fact that the action is in his option set is sufficient to its being feasible in the binary sense, and what matters for scalar feasibility is only, his having chosen that action, how likely the outcome is to obtain. An action can be really hard to do, but still feasible. Correspondingly, an action might be really easy to do but the outcome it connects to infeasible. Imagine that all I have to do to double my money is go to the casino and put all my money on black. Going to the casino is really easy, and so is putting all my money on black. But the chance of my doubling my money is slightly less than half (the odds are in the house's favour). So we might judge that the action of putting all my money on black is a less feasible way to produce the outcome of doubling my money than many other salient options. We shouldn't confuse easy and hard with feasible and infeasible. We can admit that an action is really hard for an agent but still insist (a) that it is feasible for him, in the binary sense that it is one of his options, and (b) that what is interesting for scalar feasibility is only the extent to which that action is likely to produce the outcome. It is relevant to what the agent should actually do that one of his actions would be really hard for him, but choice-worthiness and feasibility come apart, which is why feasibility supplements decision theory without being able to supplant it.

To summarize, a proposal is more feasible the less it clashes with the relevant soft constraints. These constraints are all a part of the context in which an agent – individual or collective – acts: economic, institutional, cultural and religious constraints (including positive morality), and the constraints posed by other agents, such as their motivation and psychology. Only very strong pathologies make an action less feasible for the agent whose projects we are interested in,[30] otherwise any action in their option set is up for grabs.

Now we're in a position to give the scalar version of (d):

---

30 Because in that case, even if she were to try, she would not succeed. Here 'trying' is a mental act rather than a physical act. Her pathology severely limits the efficacy of her trying.

(e) The scalar feasibility of an outcome is equal to the probability of the outcome given the best action.

And the truth conditions:

The probability of the outcome is determined by the extent to which the best action clashes with soft constraints in producing the outcome.

The logic of (e) also incorporates conditional probability. What is important is the probability of the outcome given the action – the probability that the outcome will be brought about assuming that the action is performed.[31]

VI. References

Barry, Christian & Valentini, Laura. "Egalitarian Critics of Global Egalitarianism: A Critique", *Review of International Studies*, 2008.

Brennan, Geoffrey. & Pettit, Philip. "The Feasibility Issue", Frank Jackson & Michael Smith (Eds.) *Oxford Handbook to Contemporary Philosophy*. (Oxford: Oxford University Press, 2005), 258-297.

Brennan, Geoffrey. & Southwood, Nicholas. "Feasibility in Action and Attitude", *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz.* T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson (Eds.), 2007.

Buchanan, Allen. *Justice, Legitimacy, and Self-Determination. (*Oxford: Oxford University Press, 2004).

Cohen, Gerald. *Why Not Socialism?* (Princeton: Princeton University Press, 2009).

Cohen, Gerald. "Why Not Socialism?", E. Broadbent (Ed.) *Democratic Equality. What Went Wrong?* (Toronto: University of Toronto Press, 2001), 58-78.

Collins, Stephanie. "Collectives' Duties and Collectivization Duties", *unpublished manuscript*, Feb. 2012.

Cowen, Tyler. "The Importance of Defining the Feasible Set," *Economics and Philosophy* 23 (2007), 1-14.

Estlund, David. "Human Nature and the Limits (If Any) of Political Philosophy", *Philosophy & Public*

---

31 Decision theorists have found all sorts of problems in doing things with conditionals, and they have resolved them by turning to counterfactuals with imaging (see e.g. Hájek, 2002). If the use of conditional probabilities in making feasibility assessments turns out to be beset by similar problems, one option would be to make the same move and exchange the conditional probability in the scalar definition for a counterfactual with imaging.

Affairs 39 (2011), 207-237.

Estlund, David. *Democratic Authority: A Philosophical Framework*. Princeton; Princeton University Press, 2008.

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

Gilabert, Pablo. "Feasibility and Socialism", *Journal of Political Philosophy*, forthcoming.

Gilabert, Pablo. & Lawford-Smith, Holly. "Political Feasibility: A Conceptual Exploration", *Political Studies* [early view], 2011.

Goldman, Alvin. *A Theory of Human Action*. (New Jersey: Englewood Cliffs, 1970).

Goldman, Holly. "Dated Rightness and Moral Imperfection", *The Philosophical Review* 85 (1976), 449-487.

Goldman, Holly. "Doing The Best We Can", in A. Goldman & J. Kim (Eds.) *Values and Morals* (Holland: Reidel, 1978), 185-214.

Hájek, Alan. "Counterfactual reasoning (philosophical aspects) – quantitative", N.J Smelser & P.B Baltes, (Eds.) *International Encyclopedia of the Social and Behavioural Sciences*. (Oxford: Elsevier, 2002), 2872-2874.

Hawthorn, Geoffrey. *Plausible Worlds*. (Cambridge: Cambridge University Press, 1991).

Jackson, Frank. "On the Semantics and Logic of Obligation", *Mind* 94 (1985), 177-195.

Jensen, Mark. "The Limits of Practical Possibility", *Journal of Political Philosophy* 17 (2009), 168-184.

Lawford-Smith, Holly. "Ideal Theory – A Reply to Valentini", *Journal of Political Philosophy* 18 (2010), 357-368.

Lawford-Smith, Holly. *Feasibility Constraints for Political Theories*. PhD Dissertation, Australian National University, 2010.

Miller, David. "Political Philosophy for Earthlings", David Leopold & Marc Stears (Eds.) *Political Theory: Methods and Approaches*. (Oxford: Oxford University Press, 2008), 29-48.

Pogge, Thomas. "Cohen to the Rescue!", *Ratio* XXI (2008), 454-475.

Räikkä, Juha. "The Feasibility Condition in Political Theory", *Journal of Political Philosophy* 6 (1998), 27-40.

Vranas, Peter. "I Ought Therefore I Can", *Philosophical Studies* 136 (2007), 167-216.

Wiens, David. "Feasibility in Political Philosophy: A Conceptual Framework", *unpublished manuscript*, Apr. 2012.