https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

FEASIBILITY IN OPTIMIZING ETHICS*

By Geoffrey Brennan

Justice is not the only value that calls for (appropriately balanced) implementation: other principles, sometimes competing with justice, must also be variously pursued and honored. And the facts help to decide the balance of due deference to competing principles: the facts constitute the feasible set that determines the optimal point on a set of fact-independent indifference curves whose axes display packages of different extents to which competing principles are implemented. G. A. Cohen "Facts and Principles"[1]

I. Introduction

I (and, it seems, Cohen and, for that matter, most economists) tend to conceive of normative analysis broadly as an exercise in *optimization*—as a confrontation of desirability with feasibility considerations, as we might put it. The simple thought is that feasibility considerations are no less relevant than are pure desirability considerations in the exercise of "doing the best we reasonably can" in the world that we actually inhabit.

An implication of this optimization picture is that moral philosophers and economists (social theorists more generally) *need* each other—the former to provide an appropriate specification of "desirability" and the latter to specify the feasibility set. The implied division of intellectual labor obligingly falls rather neatly along disciplinary lines; and it is perhaps tempting to conclude (comfortingly) that disciplinary specialization is just fine—that economists and philosophers can happily proceed in isolation and leave it to the smaller company of interdisciplinary brokers simply to integrate the inputs from both sides.[2] I do not think any such simple division works and one of the tasks of this essay is to indicate why.

* I am grateful to David Estlund, Nic Southwood, and David Wiens for comments on an earlier draft. I thank the contributors to this volume, especially Guido Pincione, the editors of *Social Philosophy and Policy*, and Dan Russell for excellent discussion of the essay and guidance in its improvement. Final responsibility, however, rests alas with the author.

[1] G. A. Cohen, "Facts and Principles," Philosophy and Public Affairs 31, no. 3 (2003): 211–45.

[2] I am especially interested in the "interdisciplinary" aspects of the issues here, because it seems to me that ambiguities about the role of feasibility analysis lie along the fault line between economics and philosophy.

doi:10.1017/S0265052513000150

© 2013 Social Philosophy & Policy Foundation. Printed in the USA.

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

Here is one notable problem. To begin, let us agree that the perspective from which the optimization exercise *starts* should not matter. Suppose, for example, that one begins with a specification of the basic normative "principles" and their relative significance at various levels of realization, and derives thereby the relevant (fact-independent) indifference map; then one proceeds to confront this indifference map with the prevailing feasibility set. Or suppose alternatively that one begins with the feasibility set and proceeds to check out the normative desirability of alternative feasible points. In either case, the final "optimum" should be the same. Like engineers constructing a tunnel, it ought to be possible to start out from either end (or from both ends simultaneously)—and end up arriving more or less at the same place.

Yet, as a matter of casual observation, it is obvious that this property of "starting-point-independence" does not hold. Those who begin at the feasibility end seem to be much more modest in their understandings of what is actually possible, while those at the desirability end appear to be either excessively optimistic, or in a few cases, simply uninterested in meeting up with the practitioners at all.[3] I think this is a serious problem and I believe that one aspect of the overall difficulty is the way in which philosophers tend to formulate feasibility considerations.[4]

Let me elaborate briefly. Many philosophers begin the discussion of feasibility considerations with the principle that "ought implies can." The idea here is to treat feasibility as a weaker form of "can"—to loosen the grip of the "can" proviso in ways that would be hospitable to the incorporation of "economic and political" considerations. This approach seems to encourage a view of feasibility as an independent *reason* for the endorsement of certain options—to be treated like one *aspect* of overall desirability. In my view, however, feasibility (facts) and desirability (principles) ought to be kept conceptually separate.

To be fair, I think the approach of weakening the "can" constraint gets *some* things right. It recognizes that "feasibility" comes in degrees much like normative "principles" do. But it encourages the analyst to think in terms of a "feasibility score," to be weighed against a desirability score—and that is, as I hope to show, an error. As we might put it, "ought implies can" *frames* the discussion of feasibility considerations in a particular

[3] Many of the issues relevant to starting-point dependence seem to be at stake in the fraught "debate" (if so it can be described) between "ideal" and "non-ideal" theory; that debate, however, is so mired in mutual misunderstanding that it does not seem to me to be a good point of departure in trying to come to grips with the issues.

[4] This is not the only problem. Economists routinely work with a model of human motivation that excludes (or at least seriously downplays) the role of ethical considerations in motivating agent behavior. But the aim of the paper is not to apportion blame as between economist/social theorists and philosophers for the failure of "starting-point independence." I take my audience here to be philosophers—and I focus my remarks on what I see as failures of many philosophers.

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

way; and we can think of the error to which that approach gives rise as a framing error.

Part of the problem may lie in the term "feasibility" itself, so I want at the outset to emphasize one feature of the "facts/principles" dichotomy as Cohen employs it: namely, that "feasibility considerations" are supposed to incorporate *any and all* facts about the world that are relevant to normative analysis. So, if only *some* of the relevant facts are caught up in a treatment of "feasibility issues," or if an analysis of feasibility tends to foreground certain facts and background others, then so much the worse for an approach in which feasibility as such is the primary focus.

Part of the problem also lies in a lack of clarity about what the object of evaluation is: whether an individual action, or a policy, or an institutional arrangement, or possible states of the world to which policy arrangements or institutions or actions might give rise.

My central aim in what follows is to show how feasibility and desirability considerations ought to be put together. In the process, I want to emphasize the demands that a normative approach that properly recognizes feasibility considerations makes on the formulation of desirability, on the formulation of feasibility, and on how the two elements fit together.

The argument is developed in the following way. In Section II, I discuss briefly what is involved in the general optimization approach, focusing on a couple of features that are relevant to the subsequent argument. In Section III, I direct attention to the question of what it is that is to be evaluated. Section IV emphasizes the inherent "fuzziness" of the feasibility set. Section V sketches out the line of reasoning that begins with "ought implies can" and ends with a feasibility score; the section also indicates why this approach can lead to erroneous choices. Section VI offers a brief conclusion.

II. Ethics as Optimization

The optimization approach to normative analysis, which I think is more or less characteristic of the "economic way of thinking," involves a rather simple-minded extrapolation from the "rational actor" model of market choice. The latter model imagines an agent with a given set of preferences over goods and a fixed amount of money, operating in a supermarket. The "feasible set" in this setting is composed of all possible bundles of the goods that she can purchase with her given amount of money. Her preferences can be variously represented—but the essential structural features of the model involve the agent doing the best she can, given the constraints imposed by her limited budget.

Those structural features are replicated in the normative optimization analog. Instead of the agent's actual preferences, we might think of Cohen's

"indifference map" as showing the preferences that she *ought* to have—or just more generally, combinations of levels of achievement of different normative principles that involve equal amounts of overall normative desirability.[5] Although the dimensions over which the optimization exercise is carried out are different and the nature of the feasibility constraints much more complicated, the optimization structure remains the same—and the indispensability of both "desirability" and "feasibility" elements remains crucial.

There are many features of this normative optimization exercise that are worthy of comment; but here I want to focus on just two: the intrinsically "comparative" aspect of normative principles, and "action guidingness" as the normative ambition.

A. "Comparative principles"

In any normative optimization exercise, it is crucial that one be able to explicate each principle in terms of degrees of realization—in terms, that is, of more and less. It is, in particular, inadequate to specify the principle in terms of requirements that have to be satisfied or ideal levels to which one might aspire. So, for example, the requirement in relation to "justice" would be not so much whether justice is or is not achieved in a particular case, but rather what more justice or less would amount to over the relevant range. Perhaps in the case of some principles (or dimensions of desirability), there will be a specific extent of achievement that is normatively salient (the "ideal" level, or the "required" level or some such); but the correct way within the optimization framework to recognize such salience is in terms of the shape of the indifference map. If the principle is most naturally described in terms of some benchmark degree of realization, then there has to be some measure of how close to the benchmark different outcomes are (in both directions) as well as the value of the increment in each case. This does not rule out cliffs and ravines in the overall "desirability" function, but the entire optimization approach tends to be undermined if an element of value is such that it has value only if realized to a certain extent (and negligible further value when it is more extensively realized). Stipulations of principles in terms of "on/off" requirements simply invite "impossibility theorems." So when Cohen makes reference to indifference maps and trade-offs, he is implicitly denying a structure of the normative space that has radical discontinuities of value.

[5] There is an issue as to whether this kind of procedure involves "reduction" of the various principles to some master value—"overall normative desirability"—or whether it simply shows the appropriate terms of trade between those different principles at various levels of realization. Ruth Chang "All Things Considered" *Philosophical Perspectives* 18 (2004): 1–22 argues the former position against positions like those of, say, Brian Barry, *Sociologists, Economists, and Democracy* (Chicago: University of Chicago Press, 1970).

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

Economists frequently complain that much of the work that moral and political philosophers do does not satisfy this "degrees of realization" demand. As John Broome remarks[6] in a related connection: "Economists tend to think comparatively and philosophers tend to think categorically." Broome thinks that economists have the better of that comparison. He makes this remark in terms of the claim that "betterness is fully reducible to goodness"—something that sounds like an ontological proposition. But without getting involved in ontological issues, it seems clear that, in a framework where trade-offs will typically be required, all "principles" must come with an associated metric that indicates the "degree of realization."

B. Action guidingness

When the economist engages in normative analysis, she wants to inquire what should be *done* and what should not be done. The setting envisaged is typically one in which there is a choice over alternative policies or alternative institutional arrangements to be made—and advice is sought on which option is "best" or "most right" or "least wrong." Economists just take it for granted that giving such action-guiding advice is the central purpose of normative inquiry. This is indeed so natural to them that they find it hard to imagine that normative analysis could have any other purpose.

But there could be another purpose. Normative analysis could be "attitude guiding." It could tell us what to think of someone who does a particular thing, or who makes a particular remark or thinks a particular thought. It could also simply tell us something of the *nature* of goodness or rightness without any particular context in mind. Just what is at stake in the attitude/action distinction when it comes to normative analysis strikes me as a large question—and I wish to finesse it here. But one context in which attitude guidingness and action guidingness come apart is in relation to assigning blame or praise. It can be the case that one is called on to assess a feature of some person or situation when there is no action as such to be performed. As we shall see, that distinction can be in play in relation to "ought implies can"; and so it is as well to register it here.[7]

[6] He makes the remark in his essay "Goodness is Reducible to Betterness: The Evil of Death is the Value of Life" in John Broome, *Ethics out of Economics* (Cambridge: Cambridge University Press, 1999), 162–74.

[7] Two other contexts in which the distinction comes into play are in relation to (i) voting because voting is much more like an expression of an attitude than it is an action in the full sense (see Geoffrey Brennan and Loren Lomasky, *Democracy and Decision* [New York: Cambridge University Press, 1993]), and (ii) the forces of esteem where one's attitudes can have *consequences* for others' actions even though they are not developed with those consequences in mind (see Geoffrey Brennan and Philip Pettit, *The Economy of Esteem* [Oxford: Oxford University Press, 2004]).

III. What are we Evaluating?

I have said both that economists conceive of normative analysis as action guiding and that they typically want to assess alternative policies or political institutions. There is a problem. Alternative policies, or alternative political institutions, even partial descriptions of states of the world, are not normally objects of choice—or at least, they can only be formulated as the outcomes of "action-guiding deliberation" if one is content to admit collective actors. Economists are very reluctant to do this in general—they tend to be methodological individualists. So strictly, they ought to treat *individual* actions as the objects of normative advice—they ought to worry about unpacking the "we" in "we Americans have to choose between X and Y." This is a problem of some magnitude, but I am not going to tackle it here. Rather, for the purposes of the exercise, I am going to allow that policies and institutions and states of the world (under partial description) *are* proper objects of evaluation—and leave to another occasion the question of whether (and in what sense) the normative assessment of them *can* truly be action guiding.

Note however that when we come to evaluate policies or institutions or partial descriptions of the world, those objects of evaluation lie somewhere along a notional path between the individual actions that can be thought of as direct objects of choice and the ultimate evaluative principles (the dimensions of desirability—the axes of Cohen's indifference map). Objects of evaluation are in this sense halfway houses between individual actions and ultimate criteria of assessment. And relevant facts about those objects are of two kinds—those that bear on the forward-looking path to the ultimate evaluative principles, and those that look back to the underlying actions. Note that both sets of facts are matters of feasibility: feasibility considerations just *are* the set of relevant facts.

An example might help to make the point clear. In 2002, George W. Bush became attracted to the idea that a significant increase in home ownership among minorities would be a very good thing. It would enable families otherwise excluded from property ownership to clamber onto the middle class gravy train. Bush was not the first person to have this idea: Bill Clinton had had it some years before. And so had many others who had noted the long-term appreciation in real estate values and thought that home ownership could be a route to diminished poverty and wider stakeholding in the American dream.[8] As Bush recognized, in creating millions of new homeowners, "*I set an ambitious goal. Achieving that goal is going to require some good policies out of Washington"!*[9]

[8] Margaret Thatcher in Britain had implemented the idea through the sale of "council houses" often at highly concessional rates.

[9] George W. Bush, Address to a Conference on Minority Homeownership at George Washington University (October 2002), http://georgewbush-whitehouse.archives.gov

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

One might pose two questions about such a goal. The first is whether it is "feasible" in the sense that the relevant Congressional/Senatorial support can be garnered—a question which presumably turns on judgments about the electoral attitudes to the policy, including the attitudes to the tax hikes necessary to make it happen (or the alternative policies that might have been pursued instead). That is a "feasibility" issue. But, second, facts bear also on the issue of whether the policies involved in pursuit of that goal would have desirable consequences overall—looking forward, as it were, to the implications of this goal for well-being and justice and liberty (or whatever Cohen's "fact-free" principles are). This forward-looking aspect also involves "feasibility" issues—but not the feasibility *of* the goal so much as the "feasibility considerations in connection with" the goal. If it seems an odd use of language to describe this latter set of facts as relating to "feasibility" at all, then perhaps the choice of terms is proving unhelpful. Economists would, I think, have no problem with the idea that feasibility considerations include the consequences of policy, no less than questions as to whether the policy can actually be implemented. Feasibility issues are those that determine the "feasibility set" in the space of ultimate evaluation.

But it is perhaps tempting for policy makers/advocates to just assume that their favorite policy is a "good thing" and to direct attention only to whether it is feasible and/or how to go about making it so. That is certainly the spirit of Bush's remarks.

In somewhat the same spirit, Martha Nussbaum lays out her blueprint under the "capabilities approach." This involves "announcing some ambitious goals"—such as "a thin system of global governance with at least some coercive powers" including, for example, a "set of world environmental regulations with enforcement mechanisms plus a tax on the industrial nations of the North to support the development of pollution controls in the South." But Nussbaum acknowledges the feasibility challenge. As she puts it:

> . . . a great deal more remains to be said about how precisely the approach can be used to generate political principles for today's world. To some extent, this job is a practical job, a job for economists, political scientists, diplomats and policy-makers. Philosophy is good at normative reasoning and at laying out general structures of thought. . . . [But] any very concrete prescriptions for implementation need to be made in partnership with other disciplines.[10]

/news/releases/2002/10/20021015-7.html. I am grateful to Dan Russell for this and the ensuing reference.

[10] Martha Nussbaum "Beyond the Social Contract: Capabilities and Global Justice" in *The Political Philosophy of Cosmopolitanism,* ed. G. Brock and H. Brighouse (Cambridge: Cambridge University Press, 2005), 211.

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

On its face, this seems hospitable to feasibility analysis. But the way in which feasibility considerations are actually admitted seems distinctly Bush-like. The role of the philosopher is to set the broad policy/institutional goals, and the role of the economist and other practical types is to make those goals happen. My point here is just that whether the "ambitious goals" are in fact desirable when their consequences for ultimate desirability principles are traced through is every bit as important a question as whether the policy/institutional goal is likely to be implemented; and both questions invoke "feasibility considerations."

To summarize, when we are dealing with policy goals that lie somewhere between the actions that realize them and the desirability principles that justify them, we need to recognize two kinds of applications of feasibility analysis: what we might call "feasibility of" applications and "consequences of" applications. Both applications invoke facts relevant to the constitution of the feasible set, as Cohen specifies it. In that sense, we might describe analysis that pursues the consequences of policy (or of institutional change) as "feasibility *analysis with respect to*" to distinguish it from "feasibility *of*"—just to remind ourselves that feasibility questions are not exhausted by "feasibility of" questions.

IV. Feasibility and Fuzziness

I referred earlier to John Broome's distinction between comparative and categorical thinking; and his suggestion that the economics/philosophy divide correlates with that distinction. I emphasized in that connection the necessity of specifying normative "principles"—the dimensions of ultimate evaluation—in terms of "more and less." Greater attention needs to be given, in my opinion, to developing in each case a plausible metric of the extent to which the various "competing principles" are realized.

I now want to make a similar point about "feasibility"—or more accurately about the nature of the "social facts" that constitute feasibility considerations. On the face of things, when we ask whether an object of evaluation is feasible, we are appealing to a concept that is naturally categorical: X is either feasible or it is not. And in many of the more standard economic applications (assessing alternative packages of consumer goods when shopping with a fixed total budget) that may be so. But in the extrapolation to the normative domain, the picture is more complicated. In general, we simply do not—and probably cannot—properly hold propositions about the feasibility of any object of assessment with total certainty. The things we know about the social world are just not of this type. What we need instead is something like a *metric* of feasibility. Feasibility, like desirability, comes in degrees of realization.

In part, the issue here is epistemic. We simply do not have the relevant evidence on many issues that would be required in order to establish whether something is feasible or not. In part, however, fuzziness is intrin-

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

sic to social context: fuzziness reflects the fact that social phenomena emerge as the aggregation of actions undertaken by large numbers of (somewhat) independent actors. Take, for example, the proposition that demand curves are downward sloping. This claim leads us to predict that when the (marginal) cost of something traded goes up, less of it will be bought. But of course, this prediction is not an *analytic* truth: individuals' preferences for the good may increase for independent reasons. Or the price increase may be interpreted as a signal that the price will increase even more in the future and induce stockpiling now. Economists can confidently *predict* that quantities bought will diminish. Equally, they can treat the increase in the price of gasoline in the 1970's as an *explanation* of the shift to smaller, more economical cars. But the claim about demand curves is a general one that might turn out to be wrong in any particular case. If an economist is advising the boss on the company's pricing policies, the prediction that a price increase will reduce sales below what they would otherwise be is a prediction that can be made with confidence. But social life is full of surprises. It is not possible to be totally certain.

Of course, to say that social scientists do not know everything is not to suggest that they know nothing. It is rather to observe that propositions in the social sciences ought properly to be hedged with provisos. Different propositions should be understood "probabilistically," with the probabilities varying according to the weight of evidence and the nature of the claim.

But what probabilities are relevant? That is a question that we must now address head on.

V. The Feasibility of, and the Taint of "Ought Implies Can"

It is tempting to begin conceptual analysis of feasibility with "feasibility of" in mind; and it is no less tempting to take, as a point of departure, the "ought implies can" (OIC) restriction. I want to show in what follows that OIC is a *bad* point of departure for discussion of feasibility considerations, even if for many moral and political philosophers it is an attractive one. After all, virtually every moral and political philosopher accepts OIC as something to be taken seriously, even if some would want to modify it in some way in some cases.

So let me begin with a remark about OIC. Recall the distinction laid out earlier between action-guiding and attitude-guiding aspects of the normative exercise. Often (mostly, perhaps) OIC works as an excuse giver. The relevant form of OIC is of this kind: if you could not φ then you cannot be blamed for not φ-ing. "Cannot" releases you from "ought." This is, as I see it, an attitude question: it asks what is our proper attitude to someone who failed to do something that was, on its face, more desirable normatively than what she actually did. Whether, and under what

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

circumstances OIC succeeds in this excuse-giving task, is a matter of debate to which I do not want to contribute here. My sole point is that if we see OIC instead as action guiding—as giving advice when the issue at stake is "what to do"[11]—then not doing what you cannot do seems pretty unexceptionable.[12,13]

If we understand feasibility in the OIC spirit, it may perhaps be tempting to define feasibility in terms of:

> F1: *X is feasible (for actor A) if there exists some action φ such that the probability of achieving X if A φ's is greater than zero.*

$$[X \text{ is F1 for A} \Rightarrow \exists\, \varphi \text{ s.t. } p(X|\varphi) > 0]$$

Simply put, if there is any chance at all that X could be brought about, then X is to be judged F1. But given what I observed in the previous section, adopting F1 means that the kinds of feasibility considerations that are specific to the social sciences can virtually never rule out anything as infeasible. Unless X is ruled out as a matter of logic (or perhaps the laws of physics) the probability that X is F1 will be close to 1. And this in turn will mean that the moral/political philosopher can proceed to rank alternatives simply on desirability grounds: understanding feasibility as F1 is pretty much license to dispense with social facts altogether.

So, to the extent that "taking feasibility seriously" involves taking account of the kinds of considerations that social science incorporates, we must reject F1. What F1 allows is a drifting into the very remotest reaches of the fuzzy feasibility set—against all the odds, as we might put it. For example, winning the lottery, given that you buy a ticket, is F1-feasible. And since buying a ticket is also feasible, then application of F1 says that the only question you need ask is whether winning the lottery is desirable. I take it as self-evident that winning the lottery *is* desirable; and that it *is* F1-feasible; but I take it as equally self-evident that these facts do not show that buying a ticket is desirable! We must reject F1.

What we need, one might think, is to set the bar at a proper height. So consider:

> F2: *X is feasible (for actor A) if there exists some action φ* such that the probability of achieving X given that A φ's lies above some threshold value p* (very much greater than zero).

[11] As distinct perhaps from what one *ought* to do. Nicholas Southwood has pressed this distinction on me. See his "'The Thing to Do' Implies 'Can'" *Nous* (forthcoming)

[12] This does not establish that one ought not try to do what one cannot do—though what one's best attempt would be in such circumstances deserves to be spelled out.

[13] The distinction between what one ought to do and what one should feel bad about not doing seems to be a significant element in cases like Professor Procrastinate—made famous by Frank Jackson and Robert Pargetter, "Oughts, Options, and Actualism," *The Philosophical Review* 95 (1986): 233–55.

This is a definition that Pablo Gilabert and Holly Lawford-Smith[14] seem to recommend. They treat feasibility-F2 as a sorting mechanism in normative analysis. If X does not meet the F2 test, they suggest, then it should be discarded from further consideration. But any X and Z that do meet the F2 test are judged exclusively on desirability grounds.[15]

But there seems to be a further problem. The whole optimization exercise is supposed to be action guiding; but strictly speaking F2 does not hinge on action at all. We could rewrite F2 as F2′:

> F2′: *X is feasible if [. . .] the probability of achieving X lies above some threshold value p*.* [All reference to A or A's action is deleted.]

Presumably, if X is F2 feasible and X is more desirable than any alternative Y that meets the F2 test, then A should do whatever it takes to bring X about—or more accurately, act to maximize the probability that X will prevail (over Y). That is, A should choose that action φ that most increases the probability that X will emerge.

This leaves us with F3:

> F3: *The feasibility of X is the probability that X will prevail,*

and with the associated F3 rule:

> F3 rule: *For any X, Y and so on, it is possible to assign a desirability score and an F3 feasibility score. If there exists an X*, such that X* has a greater desirability score and a greater F3 score than any other such X, Y and so on, then X* is overall best, and one should choose that action that maximizes the probability that X* will emerge.*

The F3 rule has at least something going for it. Even though winning the lottery is much better than retaining your ticket price, the probability that you will win the lottery is very low—sufficiently low to be F3 infeasible.

But we should reject the F3 rule. We should reject it because it omits feasibility-relevant information that is crucial in determining whether or not X* is optimal. We must be concerned not only with the probability that X, Y and so on will prevail but also with what happens when X, Y etc do *not* prevail. If X prevails with probability p(X), then X does not prevail

[14] In Pablo Gilabert and Holly Lawford-Smith, "Political Feasibility: A Conceptual Exploration," *Political Studies* 60 (2012): 809–25.

[15] It is worth noting that once we jettison F1 in favor of something more demanding it follows that lots of things that actually happen are not "feasible" in the F2 sense. I am grateful to Nicholas Southwood and David Wiens for pressing this point. On the other hand, things that "actually happen" by virtue of perhaps random movement might become F2 feasible once they have occurred: it is often quite F2 feasible to hold onto an emergent feature of the world when it would not have been F2 feasible to bring it about ex ante. Once you have won the lottery, you can normally hold onto the winnings.

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

TABLE 1. *States of the World.*

Actions	α p(α) = 1/3	β p(β) = 1/3	χ p(χ) = 1/3
a	X	X	∼X
b	Y	∼Y	∼Y

with probability (1 – p(X))—and the ∼X outcome may be very bad. The corresponding ∼Y outcome may be much less bad; so that even if Y is not as desirable as X, and not as F3 feasible, it may be better to undertake the Y-producing action than the X-producing one.

A simple example is sufficient to illustrate the point. Suppose that there are three states of the world possible: α, β, χ and that these are equiprobable. There are two objects of appraisal X and Y; and X is more desirable than Y (and X is more desirable than ∼X; and Y more desirable than ∼Y). Further, X prevails in states of the world α and β, but not in χ; while Y prevails in state α only. So X is more F3 feasible than Y. The action that produces X is a; that which produces Y is b.

The situation is summarized in Table 1.

Let D show the overall desirability score in terms of basic principles and suppose:

$$D(X) = 100;\ D(\sim X) = -100$$

$$D(Y) = 90;\ D(\sim Y) = 50$$

Now, if we take a simple "expected desirability" calculation,[16] it is clear that:

$$D(a) = 33;\ \text{and}\ D(b) = 63$$

It is better, taking feasibility considerations fully into account, to do action b, and thereby to go for Y (and ∼Y) rather than X (and ∼X). This does not show that Y (and hence b) is optimal since there may be another outcome Z associated with action c that is better in expected terms than Y. But it is sufficient to show that b is better than a, despite the fact that X scores higher on both F and D metrics than Y does.

To put a little flesh on the abstract bones of this example, consider the issue of achieving the *best* monetary policy (X) versus the kind of monetary policy that at its best would arise under democratic political process

[16] Once we enter a probabilistic world, there is an issue about whether we should be risk neutral about desirability. That is a significant question that I cannot hope to engage here. I would simply note that risk neutrality (in relation to desirability) is a quite specific assumption and I have done nothing here to justify it.

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

TABLE 2. *States of the World.*

Actions	α bankers good politicians good	β bankers good politicans bad	χ bankers bad politicians bad
Independent central bank	Best monetary policy	Best monetary policy	Disastrous monetary policy
Political monetary policy	Mediocre monetary policy	Poor monetary policy	Poor monetary policy

(Y). Because of electoral pressures, Y tends to fall short of X even when politicians are well meaning and competent. X is more likely to be achieved by an "independent central bank" because the central bankers are relieved of electoral pressures: that is the *point* of independence! We can characterize three possible states of the world by the possible motives of central bankers and politicians respectively: each can be self-serving or publicly interested. In state α both are publicly interested; in state β, central bankers are publicly interested but politicians are self-serving; in state χ, all are self-serving. Suppose there are reasons to think central bankers are less likely to be self-serving than politicians: central bankers are selected for their professional competence and their public spirit, whereas politicians self-select for their ambition and love of power, and are elected on the basis of their rhetorical gifts and natural charisma.[17] But politicians are constrained by an admittedly imperfect electoral process; so if they do really badly, they will be called to account. Self-serving central bankers can do much more damage in choosing monetary policies than self-serving politicians will be able to do: that is the downside of independence.

Then we get Table 2, conceived as a specific instance of Table 1. The simple point of this example is that it will be better, all things considered, to opt for democratic determination of monetary policy than to opt for an independent central bank even though better monetary policy is more desirable and more F3 feasible than leaving monetary policy in the hands of the politicians. (I offer this example as an illustration. I do not commit to the claims that central bankers are less likely to be venal—or motivated in ways contrary to the public interest—than politicians are, or that the relevant states of the world in the three columns are necessarily equally likely. But I do not consider the assumptions hopelessly implausible. In any event, the bottom line of these examples is conceptual: we must dispense with F3 and its associated F3 rule.)

[17] The role of such factors is a theme in Brennan and Lomasky, *Democracy and Decision* (New York: Cambridge University Press, 1993).

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

It may seem as if I have engaged in a long process of successive elimination of notions of "feasibility of" when it would have been simpler to offer my own preferred interpretation. I *will* suggest what I regard as a reasonable procedure shortly. But I have ordered the discussion as I have for a purpose—namely, to elaborate on a thought that OIC is a misleading and unhelpful place to begin to think about how feasibility considerations might best be incorporated into normative analysis.

As to my preferred procedure: Recall that the object of the exercise is taken to be action guiding. That is, our ultimate objective is to scan the set of a's, b's and so forth in order to determine what it is best to *do.* So, my inclination is to take the set of underlying *actions* as a point of departure, and trace out for each such action the probability distribution of possible outcomes and the desirability score for each with respect to each of the "fundamental principles" (justice, well-being, liberty and so forth). The "feasibility" dimension of this exercise revolves around: first, specifying carefully the objects over which the chooser does have complete control (the set of actions); and second, determining the outcomes and their associated probabilities of emerging that are associated with each such action. The desirability dimension involves assessing the various outcomes with respect to the fundamental principles. This means in particular developing a plausible metric for each principle that will measure "the extent to which that principle is realized or honored," and indicating the relative values of the different principles at different points in the relevant achievement space (i.e., tracing out Cohen's "fact free" "indifference map").

This kind of exercise is demanding: it leaves plenty of work for both moral/political philosophers and social scientists. And it is tempting to look for various kinds of simplifications and shortcuts. Some of these simplifications and shortcuts could perhaps be thought of as "optimal responses to epistemic constraints." But equally some of them will involve abstracting from what is central. Two ways of abstracting in particular seem relevant: one (the economist's temptation) is to abstract excessively from desirability issues; the other (the philosopher's temptation) is to abstract excessively from feasibility considerations. The anxiety that this paper is designed to elaborate is that an approach to feasibility questions via the ought-implies-can principle is likely to be of that latter kind.

The way I believe we ought to think about "feasibility" is in terms of what is more or less plausible. Indeed, I rather wish that "plausibility" rather than "feasibility" were a term in wider currency: "plausibility analysis" would give the right kind of impression of imprecision. But "plausibility" still carries a connotation of an earlier predominant meaning—as something that looks right that is in fact wrong. And needless to say, that connotation is not one I would want feasibility analysis to carry!

It may be useful, finally, to draw out one implication of my preferred picture, because it relates to an issue over which there seems to me to be much confusion. When someone remarks of a policy proposal or an insti-

tutional change that it is "just too costly," that remark is often met by the response that any such observation has nothing to do with feasibility: "To say that X is too costly" critics respond, "is not to say that it is infeasible."

This response is surely valid on any defensible notion of feasibility. But it is not the end of the conversation. Any statement about "costs" invokes *some* claim about feasibility. If X is "too costly" it is by virtue of some consequences of X for things that are worth consideration. If X is costly at all, it is because something else (Y or Z) is forgone; and if X is *too* costly, the claim is that some such Y or Z is worth more than X is, in terms of the realization of ultimate desirability principles. To be sure, the proper currency of "worth" here must refer to the ultimate "principles"; but the claim that X entails forgoing Y and Z is a *feasibility* claim. Or better put, perhaps, the claim is that having both X and Y is very implausible, and that having X makes having Y much more implausible. Claims about costs are claims that invoke both desirability and feasibility considerations in equal measure. To say that X is costly does not imply that X is not worth the cost. But normative optimization is an exercise in comparison; and canvassing the cost of alternative actions is an intrinsic part of that comparison. It invokes both desirability and feasibility elements in an appropriate amalgam. And that is entirely consistent with my picture of how things ought to go: desirability and feasibility elements belong together. Principles and facts may be conceptually separable—as Cohen insists (properly, I believe). But they are indispensible co-partners in any attempt to work out what it is best to do!

VI. Conclusion

"Doing the best we can" is a matter of optimization: it involves a confrontation of pure desirability considerations with feasibility considerations. In that exercise, feasibility and pure desirability are—or ought to be—equal partners. And it ought not matter for the location of the optimum whether analysis begins at the pure desirability or the feasibility end: the optimum itself ought to be "starting-point independent." But economists frequently complain that philosophers do not take feasibility considerations sufficiently seriously. And philosophers tend to think that economists do not take pure desirability considerations sufficiently seriously. I think both are right. But here, I have been concerned to address the philosophers. My intended audience is not those philosophers who think (if any do) that feasibility considerations are simply irrelevant, so much as those who acknowledge the relevance of feasibility considerations in principle but do things in practice that treat such considerations inappropriately.

As I see it there are five main messages:

1. In any optimization exercise, basic normative principles must be specified in terms of a metric of realization—in terms of "better-

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press

ness" rather than "goodness" (or greater or lesser justice, liberty, well-being, and so on).

2. Feasibility itself "comes in degrees": it is irretrievably probabilistic. Put another way, the feasible set is fuzzy. This does not mean that charges that a particular evaluand is "infeasible" are meaningless, but such claims need careful interpretation. The term "feasibility" understood as a category is probably unhelpful here.
3. The principle that "ought implies can" is a bad point of departure for understanding feasibility analysis: it frames the concept of feasibility in a misleading way. In particular, it encourages conclusions like the following:

   > if X and Y are objects of evaluation (say policies or institutional changes) and it is the case that: X is more desirable than Y; and X is more feasible than Y; then X should be chosen over Y.

   This conclusion does not follow.
4. Part of the difficulty in the optimization exercise as it often presents itself relates to the nature of the evaluands. For in working out what it is best to do, the actual options are the underlying *actions;* and what we are required to do is to trace out the full consequences of those actions for the achievement of the basic evaluative principles. This "tracing out" will typically require an assessment of the whole range of possible outcomes weighted by their respective probabilities. The fact that a good outcome arises from an action with a reasonable probability is not sufficient to make that action the best to do—we have to attend as well to what happens when things do not work out so well.
5. "Doing the best we can" is, whatever else, a matter of *doing.* There may be other tasks to which normative analysis can be directed. Such other tasks have not been my concern here. My only suggestion in this connection is that a normative analysis directed at those other tasks should not be allowed to leak inappropriately into the task of determining ethically optimal action.

Philosophy, Australian National University
Philosophy, Politics, and Economics, Duke University and University of North Carolina

https://doi.org/10.1017/S0265052513000150 Published online by Cambridge University Press