

draft as at 26 Mar 10

from Joseph Millum and Ezekiel Emanuel, eds, *Global Justice & Bioethics* (New York: Oxford University Press, 2012), pp. 153-77.

The Bioethics of Second Best*

Philosophy Program, Research School of Social Sciences
Australian National University
<goodinb@coombs.anu.edu.au>
and
Department of Bioethics, Clinical Center
National Institutes of Health
<goodinre@cc.nih.gov>

* Earlier versions of this chapter have been presented to seminars of the Bioethics Department at the Clinical Center of the National Institutes of Health and of the Georgetown-Johns Hopkins Greenwall Fellows. I am grateful to participants in those discussions, to Joe Millum for written comments and to Marion Danis, Christine Grady, Chiara Lepora, Frank Miller and Annette Rid for resource materials. This research was supported by the Intramural Research Program of the National Institutes of Health during my time as a visiting scholar in the Clinical Center's Bioethics Department. The views expressed are my own and do not represent the opinions or policies of the National Institutes of Health or the Department of Health and Human Services.

The real world is messy in all sorts of ways. Those who try to model the world, whether as scientists trying to explain it or as moralists trying to change it, attempt to abstract from that messiness. They seek models that are simpler, cleaner, more transparent than the reality that those models are attempting to mirror.¹

Newtonian physics envisages balls colliding on a frictionless plane. Of course, when Fast Eddy shoots pool on a real table, he had better remember about friction. But starting with an idealized Newtonian model and then factoring in friction serves as a pretty good guide.

Something analogous is generally supposed to be true when it comes to moral philosophy. In terminology owing to Rawls (capturing an idea that is much older) moral philosophers distinguish between "ideal theory" and "non-ideal theory."² They accept the need to make adaptations to ideal theory when applying it to the real world, of course, adjusting for the ways in which the actual differs from the ideal.³ But it is generally assumed that making those sorts of adjustments will be no more problematic for moral philosophers than it is for Fast Eddy hunched over the green felt.

¹ Black 1962.

² Rawls 1971, pp. 245 ff. Sreenivasan *infra*. Simmons 2010. For an excellent application to another realm, that of international migration and global justice, see Carens (1996).

³ Of course, Rawls does not abstract from absolutely all non-ideal aspects of the actual world. He does not abstract from the sad fact of material scarcity for example: if everyone could have as much as they wanted and questions of distributive justice would not arise (Rawls 1971, sec. 22, pp. 126-30). What grounds there are from abstracting out some non-ideal facts but not others for purposes of moral theory is a large topic, too little discussed. But I simply note it here in passing.

Starting with the ideal may indeed be ideal, in all sorts of ways.⁴ But in one important respect it might be seriously in error. It is wrong simply to assume – as moral philosophers typically do, without much elaboration or discussion – that, when it comes to applying their moral theories, starting from the ideal and tacking back to the real will be relatively straightforward. They dismiss all that with a cavalier "mutatis mutandis." As this chapter will show, it is not that easy.

From basic economic theory we know that the second-best state of affairs might be very different indeed from the first-best. Compensating variations in several dimensions might be – and typically are – required to make up for shortfalls in others. Where that is so, the right thing to do in some non-ideal world cannot be simply and straightforwardly read off ideal theory's prescription of what to do under ideal circumstances.

After briefly cataloguing various different respects in which the real world might deviate from the that presupposed in ideal theorizing, I shall introduce the General Theory of Second-best and explain the trouble it makes for reading real-world prescriptions directly off ideal-theory pronouncements. The upshot of that discussion is that adjusting our (ideal) theory to the (real, messy) world may be genuinely problematic. Alternatively, we might try instead to adjust our (real, messy) world so that it better fits the conditions presupposed by ideal theory. If we succeeded

⁴ Ideal theory, Rawls (1971, pp. 8-9; 391) says, is the more "fundamental." We need it to tell us what we should be aiming at through social reform of non-ideal circumstances (Rawls 1971, pp. 8-9, 245; 1993, p. 285; 2001, p. 13).

completely in that, then the prescriptions of ideal theory would be directly applicable, without the sort of adjustments that the Theory of Second-best makes so fraught; but unless we can completely instantiate all the conditions presupposed by ideal theory, we might still fall afoul the Theory of Second-best. Yet another response to the problems posed by the General Theory of Second-best would be to abandon the quest for excessively idealized moral theories, and instead to theorize the world in the vicinity of where we actually find ourselves. Depending upon the empirical facts of the matter, any of those strategies might work. And any of them might, again depending on empirical facts of the matter, be usefully supplemented by a pair of partial solutions, decomposing problems or making policy choices that are relatively robust across changing circumstances. In the end, however, there is no sure solution to the problems posed by the Theory of Second-best – only a suite of more or less imperfect alternatives.

I. Why Second Best?

The need for “second-best” solutions arises because “first-best” solutions are unavailable. Moral philosophers mark this as a contrast between “ideal” and “non-ideal theory,” and it is that moral usage with which this chapter is concerned. (There may be various other non-moralized senses in which a

state of affairs can be non-ideal. My focus here, however, is on deviations from the ideal that are problematic in a more principled way.⁵⁾

At a minimum, we can say that state of affairs is morally non-ideal if resources are inadequate to meet morally obligatory tasks.⁶ This shortfall might occur in any (or several or all) of the following dimensions:

- 1) Material resources might not be adequate. Morally, we ought to prevent all preventable deaths, let us suppose. But given available technology, we simply cannot manufacture enough vaccine in time. Or if health dollars are strictly limited, we might decide to settle for a second-best procedure that is "much cheaper and almost as good" for treating one condition, in order to free up health dollars to treat some other condition.⁷ Or poor country that relies on foreign donors with short time horizons might opt for a narrowly targeted health intervention that shows quick results rather than investing in more comprehensive primary health care infrastructure: the latter would be better, but the foreign funds would not be available for it.⁸

⁵ Whereas Hume and Rawls (1971, pp. 126 ff.) following him would define the "circumstances of justice" in terms of scarcity of resources relative to desires, I shall define what's morally "ideal" or "non-ideal" relative to moral rights and duties.

⁶ I am uncertain how to incorporate imperfect duties and supererogation here. Maybe it morally matters that people have enough resources to perform some of those acts too, and a state of affairs is also morally non-ideal insofar as they do not.

⁷ The possibility might be more theoretical than real, judging from the scarcity of reports of such "decrementally cost-effective medical innovations" in the published literature. A survey of medical cost-utility analyses published between 2002 and 2007 found that only 9 of the 2128 interventions described were \$ 100,000 or more cheaper than the existing standard of care per Quality Adjusted Life Year sacrificed (Nelson et al. 2009). It is an open question whether this result is due to a genuine paucity of "much cheaper and almost as good" interventions or whether it is merely an artifact of publication practices.

⁸ Msuya 2003, p. 17.

- 2) Ideational resources might not be adequate. We might have inadequate or even erroneous understanding, either normatively (of True Morality) or empirically (of how the world works). We might not know how to make a vaccine. Or we might not understand why morally we should make and distribute it widely, even if we could.
- 3) Institutional resources might not be adequate. Morally, let us suppose, we ought to harvest organs from willing donors immediately upon "death" and transplant them promptly into the most appropriate recipients. But in the absence of institutional arrangements for wide-scale tissue-matching and priority-setting, we often end up giving the organ to someone who, while appropriate, is almost certainly not the most appropriate.⁹
- 4) Motivational resources might not be adequate.¹⁰ Morally, we ought to allocate scarce medical resources to the people who need them most badly. But we just cannot bring ourselves to give the last available dose of vaccine to some stranger who is at greater risk rather than giving it to our own child, or to send scarce vaccine to people who are at greater risk abroad rather than allocating it to our fellow countryfolk.
- 5) Coordination resources might not be adequate. Morally, we ought to prevent the spread of life-threatening diseases, let us suppose. But even

⁹ Healey 2000; 2006. I employ scare quotes, because brain-dead patients are in other important senses not really dead: indeed, dead organs can save no one's life.

¹⁰ This the case most often discussed in the philosophical literature, by Rawls (1971, p. 8, 254, 351) when he defines "ideal theory" in terms of "strict compliance" and by others who talk of the "demandingness" of morality (Murphy 2000; Mulgan 2001; cf. Goodin 2009).

though the requisite motivational, material, ideational and institutional resources are all in place for doing so, we might still be thwarted by a failure to coordinate our well-meaning efforts. The correct sequencing of pharmacological, psychological and social interventions is crucial, let us suppose; different people necessarily have to be responsible for each; and coordinating their concerted action is beyond our powers.¹¹

The source of the shortfalls that make a situation less than ideal might be in our own resources or in the resources of others. Thus, for example, we may be unable to inoculate everyone in time because of a lack of material resources to manufacture enough vaccine ourselves on our part, or on the part of others (e.g. foreign manufacturers). Or the reason we are unable to get the vaccine to those most in need might have to do with a failure of other people's motivations rather than our own (we declined to give the last dose of vaccine to our own child, only to watch someone else appropriate it for her own child who was not particularly needy, either). And so on.

In terms of responses, it is an open question what morally we ought to do when we find ourselves in situations where resources are inadequate in any of these ways to meet morally obligatory tasks. One alternative would be to take the world as we find it and do the best we can in those (admittedly, non-ideal) circumstances. Another alternative would be to try to transform the situation – make the non-ideal circumstances more nearly ideal, wherever

¹¹ Often, but perhaps not always, institutional arrangements are the solution to such problems – in which case the fourth case collapses into the second.

we can. Both strategies risk falling systematically afoul of the General Theory of Second-best, as I shall go on to argue in Section III.

II. Second-best Might Be Completely Different

Whether or not we should "settle for less than the best" often provokes heated disputes – and rightly so. (I shall say more of that shortly.) What is involved in settling for second-best is, in contrast, typically taken to be relatively straightforward.¹² Wrongly so, I now want to argue.

Second-best solutions can, and often do, display peculiar features that pose particular challenges of institutional design and policy choice. While these are nowise unique to bioethics, neither is bioethics in any way immune to them.

A. The General Theory of the Second-best

Long ago, economists Richard Lipsey and Kelvin Lancaster proved the General Theory of Second Best. Put into what passes among economists as plain English: “The general theorem of the second best states that if one of the conditions [characterizing the optimal outcome] cannot be fulfilled a

¹² Or third-best, if what would be second-best is also unattainable, and so on: those further iterations will become important in my discussion below; but for now let me speak as if “second-best” embraced all those non-first-best possibilities.

second best ... situation is achieved only by departing from all other optimum conditions."¹³

Lipsey and Lancaster's formal proof is tied to specifics of the standard economic set-up (well-ordered preferences, general equilibrium, and so on). And their very strong conclusion ("only by") derives from some of the very particular stipulations that economists conventionally make concerning preference functions.

There is, however, a weaker – and more genuinely general – version of the Theory of Second Best that is independent of any such assumptions. In that weaker and more general form which will be my focus in this chapter, the Theory of Second Best says this: if the first-best state of affairs cannot be obtained, the second-best state of affairs is not necessarily identical to the first-best in any respect. Whereas Lipsey and Lancaster's stronger version would say "necessarily not," the weaker and more general version that I shall be discussing says merely "not necessarily."

The phenomenon is a familiar one across broad swathes of life, once you come to think about it. Here is one homely example. Suppose that my first-best car has three attributes: it is a (1) new (2) silver (3) Rolls Royce. But

¹³ Lipsey and Lancaster 1956, p. 12 elaborate: "The general theorem for the second best optimum states that if there is introduced into a general equilibrium system a constraint which prevents the attainment of one of the Paretian conditions, the other Paretian conditions, although still attainable, are, in general, no longer desirable. In other words, given that one of the Paretian optimum conditions cannot be fulfilled, then an optimum situation can be achieved only by departing from all the other Paretian conditions. The optimum situation finally attained may be termed a second best optimum because it is achieved subject to a constraint which, by definition, prevents the attainment of a Paretian optimum. " (Lipsey and Lancaster 1956, p. 11).

suppose that there is no such car available at the moment, and for some reason I really must acquire a car immediately. Hence I have to settle for second-best. The Theory of Second Best cautions me that my second-best car will not necessarily be one that displays more rather than fewer of those same attributes as my first-best car. Thus, for example, the second-best from my point of view would probably be a (1) week-old (2) black (3) Jaguar rather than a (1) new (2) silver (3) Toyota, if those were the only two cars on offer. That is true, even though the Jaguar displays none of the same features as my first-best car, and even though the Toyota displays two out of the three. That is precisely the point of the Theory of Second Best.

Avishai Margalit offers another example, this one drawn from the teachings of the Church of Rome:

The Catholic Church believes that being a nun is the ideal life. It is the life of perfection for women. The Catholic Church also believes that the sacrifice entailed in giving up sexuality and motherhood is such that most women cannot attain the ideal of becoming nuns. The second best for a woman is not to become a nun with a lax attitude toward the prohibition of sexuality, but instead to become a mother.¹⁴

Similar examples pervade the public sphere as well. Imagine a health promotion campaign. Were we promoting a healthy lifestyle, what characteristics would that have? The first-best lifestyle would (let us imagine) include at least the following attributes: (1) no tobacco, (2) little alcohol and (3) regular exercise. However, nobody would seriously think that the second-best lifestyle would be one that displayed perfectly two out of those

¹⁴ Margalit 2010, p. 116.

same three attributes without any regard to the other. The second-best lifestyle would definitely not be one in which you neither smoke nor drink but never get any exercise at all, either.

The same phenomenon also recurs when implementing social ideals more broadly. The best society, let us imagine, is one that is both (1) free and (2) equal. But suppose the only way to maintain perfect equality is to interfere with people's freedom in some respect (e.g., with their freedom to bequeath large sums of money to their heirs). While the first-best society is, ex hypothesi, one that maximizes freedom and maximizes equality, the second-best society is probably not one that maximizes completely either of those values, regardless of the cost to the other. Instead, the second-best society is probably one that scores pretty highly on both freedom and equality without literally maximizing either (e.g., imposing confiscatory taxes on bequests, but only above a certain sum).

That thought is sometimes expressed in terms of "value trade-offs."¹⁵ Of course, there may be some values you refuse ever to trade off for any others. Maybe some values stand in a strict hierarchical relation to others, such that any difference (no matter how small) on the top-ranked value trumps for you any difference (no matter how large) on the lower-ranked value. Much more typically, however, you would probably be prepared to

¹⁵ As in the title of Arthur Okun's book, Equality and Efficiency: The Big Tradeoff (1975).

give up a little bit of value-attainment in one dimension in exchange for a certain amount in the other.¹⁶

When the choice situation forces us to trade off one value for another, we have to decide which matters more (and by how much) in that situation. In such cases, worries can arise concerning the commensurability of values. Can they really be compared, in ways that would allow us coherently to trade more of one for less of the other?

Sometimes, however, choosing among options does not require any invidious comparisons across values. All the comparisons can sometimes be done within the same value. Cass Sunstein offers various examples of that sort in his discussion of "Health-health Trade-offs" that arise in risk regulation. For example, "Regulations designed to control the spread of AIDS and hepatitis among health care providers may increase the costs of health care, and thus make health care less widely available, and thus cost lives... A ban on carcinogens in food additives may lead consumers to use noncarcinogenic products that carry greater risks in terms of diseases other than cancer."¹⁷

Note that what you regard as second-best will always depend on the interaction between your evaluative standards (preferences, values) and your options (the feasible set over which you can effectively choose). But note well: it is not as if you change your standards when confronted with new

¹⁶ Barry 1965, pp. 4-8; cf. Rawls 1971, pp. 34-45.

¹⁷ Sunstein 1996, pp. 1535-6.

options. The same standards apply. They simply apply differently over a different feasible set.

Note too that the Theory of Second-best applies only to multiple-attribute decision problems. If there is only one criterion (or if there are multiple criteria that are hierarchically ranked in a lexical order, such that only one is in play at any given time), then it is necessarily the case that what is second-best will be whatever is as similar possible to what is first-best on that only or lexically-prior criterion.

There are two factors driving the Second-Best phenomenon. One is "suboptimization." That is the error of optimizing on only a subset of all the dimensions that are actually important to you. In so doing, you get the right result with respect to that subset of dimensions to which you are paying attention– but the wrong result with respect to the other dimensions that you are ignoring in the process. That is how you end up choosing the new silver Toyota – by fixating on two criteria that matter to you but ignoring the third (that you want a luxury car, not merely a new black one). Typically, the right thing to do all-things-considered differs from the right thing to do only some-things-considered.

The second factor driving the Second-Best phenomenon is "interaction" across those dimensions. So, for example, education interacts with health which interacts with employment: the more education people have the better able they are to make healthy lifestyle choices and the better able they are to take advantage of employment opportunities; and the healthier people are the better able they are to hold down a job. That is why policymakers need to

consider the entire suite of education-health-employment policies all together in a holistic manner, rather than just attending to them separately.¹⁸

III. Strategies for Bridging the Ideal-Real Gap

With the Theory of Second-best firmly in view, let me now return to show what trouble might make for moral philosophers trying to apply their "ideal theories" to the real world.

A. Make the Theory Fit the World

Standard practice, as I have said, is for moral philosophers to develop their theories of what should be done in "ideal" conditions, abstracting from various messy features found in the real world. When they come to apply those theories to the real world, they then simply make such adjustments to ideal theory's prescriptions as are required in light of those non-ideal facts about the real world.

Sometimes those are modest tweaks. Other times they are major bolt-ons. For an example of the latter, notice that in the world of ideal theory no

¹⁸ As Prince Edward Island tried to do, with little success, by giving a single regional authority responsibility for reallocating resources across a broad range of health and community services (Stoddart et al. 2006).

one would ever break the law and no government would enact an unjust law. Both things sometimes happen in the real world, however. To accommodate that fact, Rawls needed to bolt a theory of corrective justice and a theory of civil disobedience onto his theory of justice when applying it to the non-ideal real world.¹⁹

Even when whole new branches have to be added to ideal theory in applying it to the real world, however, writers like Rawls tend to presume that the basic structure set by ideal theory remains unchanged. Adjustments will inevitably be required at the margins, and extra bits will have to be added on the edges. But, writers like Rawls assume, the great bulk of ideal theory's prescriptions will remain the same as applied to the real world.²⁰

Furthermore, however far in the background, this assumption is not just some incidental oversight. It goes to the heart of the methodology. If systematic, thoroughgoing revisions to ideal theory's prescriptions would be regularly required in applying it to any real-world situation, then there would be no point in starting with ideal theory. If tacking back to the real world requires us to rethink everything afresh, then it is not at all clear what ideal theorizing has bought us.

¹⁹ Rawls 1971, pp. 8-9, 245-6, 575.

²⁰ True, Rawls offers what Simmons (2010, pp. 24) calls an "integrated ideal." But while the ideal itself is integrated, the procedure Rawls suggests for pursuit of that ideal in the non--ideal world is not. Simmons thinks otherwise. He (and he thinks Rawls) appreciates that "our attacks on particular, especially offensive injustices may be... understandably compelling." Yet he (and he thinks Rawls) suppose that "few devotees of 'partial justice' would be able to sustain their single-minded commitments in the face of clear evidence that their efforts were setting back or permanently blocking movement toward overall social justice." Yet a "one injustice at a time" approach is precisely what Rawls (2001, p. 13) seems to recommend, in

The Theory of Second-best, however, warns that that may be precisely what is required. There is no reason to think that the second-best (i.e., non-ideal) world is necessarily identical to the first-best (i.e., ideal) world in nearly every respect. There is no reason necessarily to think that most (or indeed any, much less most) of ideal theory's prescriptions will carry over unchanged (or even just minimally changed) in a world that is less than ideal in any respect.

So the Theory of Second-best stands in the way of any assumption that what is second-best can necessarily be read easily and straightforwardly off ideal theory's description of what is first-best. In adapting ideal theory to a non-ideal world, systematic root-and-branch changes might be required to ideal theory's prescriptions.

Or they might not. All that my weaker and more general version of the Theory of Second-best says is that the second-best does not necessarily bear any close resemblance to the first best, not that it necessarily does not. Whether it does or not depends just on boring empirical facts of the matter – facts to do largely with interactions across the evaluative domains and the centrality of any given evaluative criterion to the situation under consideration. This is what creates possibilities for more or less successfully "decomposing problems" in ways I shall discuss shortly.

The Theory of Second-best is thus better understood as cautionary rather than condemnatory. It does not say that the standard approach of first

saying that "ideal theory... should ... help to clarify the goal of reform and to identify

producing an ideal theory and then assuming that the right prescriptions real world are pretty much the same is necessarily wrong, always and everywhere. There might be some circumstances in which that would work fine. All the Theory of Second-best, in its weaker and more general form, says is that we must not assume that that is always the case. We must proceed with caution, and we must be prepared for that procedure sometimes to fail badly.

B. Make the Real World More Nearly Ideal

The standard strategy is to start from ideal theory and adjust it to fit the real world. The Theory of Second-best tells us that there may be real problems in doing so – that really major changes in ideal theory's prescriptions might be systematically required. Might those problems be avoided by taking the opposite tack, making the real world more ideal in the ways to which the prescriptions of ideal theory are tailored?

There are many reasons, independently of difficulties posed by problems of identifying the second best, for wanting to make the real world more ideal. Other things equal, it would be better if we were not lacking in the resources required to discharge our moral responsibilities. And all the

resources listed at the outset – material, ideational, institutional, motivational, coordination – can, in principle, be increased.

Just how hard or costly it will be to increase our moral resources is an open question, and one that is probably best answered differently for different classes of cases. It may be easier to effect changes in your own institutions or your own motivations than those of others, for example. That might suggest, in turn, that we ought to strive primarily to correct our own failings while by and large taking those of others as given. Or for another example, it may be easier to change institutions than motivations: that was Rousseau's thought in "taking men as they are and laws as they might be."²¹

Conceptually, it is important to appreciate that feasibility always be understood dynamically.²² Writers like Bentham rightly bemoan the way in which "the plea of impossibility offers itself at every step, in justification of injustice in all its forms."²³ It is the job of leaders, as Max Weber says at the end of his essay on "Politics as a Vocation," to make possible tomorrow what is not possible today.²⁴ True, "ought implies can." But that does not excuse us from doing what we ought to do, merely because we cannot do it just at the moment. If we can get ourselves into a position to do what we ought to do (in time still to do it), then that's what we ought to do, ceteris paribus.

²¹ Rousseau 1762. For discussions of these options with special reference to global justice debates see: Valentini 2009; Lawford-Smith 2010; and Ypi 2010.

²² Gilibert 2009. Simmons (2010, pp. 24-5) rightly emphasizes this, and supposes Rawls either does or should do similarly within the logic of his own theory.

²³ Bentham 1827, vol. 7, p. 285. Goodin 1982, ch. 7.

²⁴ Weber 1919/2004, p. 9.

Note that ceteris paribus clause well, however. It is not necessarily the case that the best thing to do, all things considered, is to transform a non-ideal situation into one that is morally ideal. It is not, anyway, if the resources which would be expended to bring about that transformation could be used in morally more desirable ways. In some non-ideal situations you can do almost as well as you could in the absolutely ideal one. Then transforming the situation from non-ideal to ideal may well cost more than morally it is worth.

Nor is it the case that making real-world circumstances more like those presupposed by our ideal theory will necessarily make it easier to surmise from ideal theory what we should do in the real world. It is not, anyway, if we do not succeed in instantiating all of the conditions presupposed by ideal theory. Only if we did could we be sure that the prescriptions of ideal theory are precisely right for us in the world that we actually occupy. If any one of the conditions presupposed by ideal theory is missing, then the Theory of Second-best warns that we might (not necessarily will, but might) need to make systematic alterations right across the board in the prescriptions of ideal theory.²⁵ Again, if ideal theory's prescriptions were themselves decomposable, and if we could make real-world circumstances completely ideal in every respect relevant to any given prescription, then in that very

²⁵ There is, furthermore, an bootstrapping issue here: if we are in a non-ideal world and we want to make it ideal, how do we know how to do that? We cannot (necessarily, anyway) use ideal theory to guide us how to get to the ideal, because *ex hypothesi* we are not in the ideal state yet; and the Theory of Second-best tells us that what is the right thing to do in the ideal state might be very different from the right thing to do in non-ideal states.

special case we could know with confidence what to do judging just from ideal theory alone. But that is contingent, a special case rather than a standard solution.

C. Theorize the Vicinity of the Actual

If the Theory of Second-best stands in the way both of fitting ideal theory to the real world and of fitting the real world to ideal theory, then perhaps we ought abandon ideal theory altogether. That is to say, instead of abstracting from all the messiness of the real world and assuming circumstances that are much more ideal than the actual circumstances in which we find ourselves, maybe we should build our moral theories from the start around assumptions that are broadly true of the world we are actually in.²⁶

Again, there is much to be said for this strategy, quite apart from any assistance it might provide in avoiding problems posed by the Theory of Second-best. Perhaps the main thing to be said for theorizing the vicinity of the actual is that we test our moral theories against our moral intuitions. Those moral intuitions, in turn, have been formed around the standard sorts of cases commonly faced in ordinary life. If we try to theorize a world too far

²⁶ For critical discussion of this option see Simmons (2010, pp. 30 ff.).

from our lived experience, we have nothing reliable to test our theories against.²⁷

While the strategy of theorizing in the vicinity of the actual is thus very tempting, there are also obvious problems with it. For one, it is messy – complicated, confusing – to try to model everything at once, without abstracting anything away. For another, theorizing that is tightly bound to existing circumstances may not give us much guidance as to the direction in which we should attempt to change those circumstances for the better, morally. For yet another, theorizing in this mode may be highly unstable, prescribing very different things as circumstances change, perhaps even only slightly.

Beyond all that, however, would this strategy really solve the problems that the Theory of Second-best poses for moral theorizing? It would, if our moral theory were crafted to the exact circumstances actually obtaining in the real world. Then our moral theory's prescriptions would be precisely right for the circumstances in which we apply it. But the Theory of Second-best warns us that, if the actual circumstances are different in any respect whatsoever from those presupposed by our moral theory, then that theory's prescriptions might need to be systematically altered. And requiring that moral theorizing be perfectly tailored to actual circumstances, when circumstances are constantly changing, would have us constantly retheorizing – leaving us with little time for acting. If for whatever reason

²⁷ Goodin 1982, ch. 1.

circumstances change faster than our moral theories can, this strategy will offer us no certain protection against the problems posed by the Theory of Second-best.

One comforting thought (which may or may not actually be true) might be this. Perhaps as a brute empirical fact of the matter the prescriptions of a moral theory would not change very much with small changes in circumstances. If that is true, then a moral theory concocted for one set of circumstances would be pretty much valid anywhere in the near neighborhood of those circumstances. If that is true as a matter of empirical fact, then a moral theory framed around our actual circumstances will provide pretty good guidance anywhere in the vicinity of our currently-actual circumstances. (And by the same token, a moral theory framed around ideal assumptions will provide pretty good guidance anywhere in close vicinity to the circumstances presupposed by that ideal theory.²⁸) If that is how things empirically turned out, then well and good. But again, that is a purely contingent empirical matter, and it could of course turn out otherwise.

IV. A Partial Solution: Decomposing Problems

²⁸ Rawls's (1971, p. 351) talk of "a state of near justice" suggests that that might have been what he was thinking, when assuming that his theory of civil disobedience could be bolted onto his theory of justice without altering much else in that larger theory.

The three grand solutions for coping with the problems posed by the Theory of Second-best all thus fail, or anyway may well fail. Let us now turn to a pair of more partial solutions that might work in a more limited way. Both turn out also to be vulnerable to threats from the Theory of Second-best, each in its own way.

Public policy-makers commonly try to decompose complex problems into their component parts, assigning responsibility for solving each component to separate actors. That might be a particular sector of society (the family unit), a particular government department (Health or Education or Employment), a particular country (the one in which the victim happens to live). In so doing, policy-makers are hoping that what is the best decision within each part, taken one-at-a-time, will when aggregated be best overall.

The basic idea obviously long predates him, but it is to Herbert Simon that we owe the modern formal concept of a "(nearly) decomposable system." Such a system is in his terms characterized by a high degree of modularity. That is to say, interactions occur (almost) wholly within modules, with (almost) none occurring across modules.²⁹

Some problems lend themselves to that sort of a solution. Others do not. It all just depends on whether we can carve up the problem into (largely) self-contained parts, and assign each part to a (largely) self-contained module. If no matter how we carve up the problem, different parts will interact too

²⁹ Simon 1969; 2002a, b.

much with other parts, then trying to divide up responsibility for the problem among (largely) self-contained units will not yield a good solution.

The Theory of Second-best however warns us that problems are not necessarily decomposable in the way that that approach requires. If something goes wrong in one dimension, then adjustments will typically be required in all other dimensions. Assigning responsibility for each dimension to different agents acting largely in isolation from one another precludes across-the-board adjustments.

Phrasing the point in economic terms, Lipsey and Lancaster write: "It should be obvious ... that the principles of the general theory of second best shows the futility of 'piecemeal welfare economics.' To apply to only a small part of an economy the welfare rules which would lead to a Paretian optimum if they were applied everywhere, may move the economy away from, not toward, a second best optimum position."³⁰

Of course, we wish it were otherwise. Nearly-decomposable systems would be quite convenient. They allow us to take advantage of specialization and of the division of labour. They admit of easy repair, replacing one malfunctioning module with another without any interruption to any other part of the system. Nearly-decomposable systems evolve and adapt more quickly to changing circumstances, and hence enjoy an evolutionary

³⁰ Lipsey and Lancaster 1956, p. 17.

advantage.³¹ So well may we wish that we were able to solve all our problems in that way.

But wishing does not make it so. Decomposable solutions presuppose decomposable problems. And where the second-best is radically different in many dimensions from the first-best, devolving responsibility for different of those dimensions to different non-interacting agents clearly risks folly.

This general principle has implications for public policy and social problem-solving across a wide range. I shall illustrate it here by reference to bioethical issues relating to health care, at three levels: individual physicians; institutional designers; and policy-makers.

A. Role-differentiated Individual Responsibilities

In addition to the general duties and responsibilities that each of us has merely as a moral agent, most of us also have “special responsibilities” toward certain other people and for certain other actions and outcomes. The role-differentiated responsibilities of health-care providers are a case in point.

The key thing to notice, in the context of the present discussion, is that the patient-provider relationship is a bilateral relationship. Each patient is assigned one or more health-care providers, who have a special responsibility for that patient's health in a way that they do not for other people who are

³¹ Simon 2002a,b.

not their patients. In the words of the oath that the World Medical Association's Declaration of Geneva asks physicians to swear: "The health of my patient will be my first consideration." Or as a philosophical lawyer (who also served Reagan's Solicitor General) elaborates: "Doctors.. owe a duty of loyalty to their clients, a loyalty which ... requires taking the medical... interests of that client more seriously than the interests of others in similar or greater need, more seriously, indeed, than formulas of either efficiency... or fairness... would require or even permit."³²

This strategy of assigning specific individuals responsibility for the health-care needs of specific other individuals is clearly a strategy of "decomposition" of the sort just discussed. It involves decomposing the problem of caring for the health of people in general into a problem of one particular person's taking care of particular patients, one-at-a-time.

That would be a good solution, if the problem were itself decomposable. But insofar as there are interactions among the health (and hence health-care needs) of different people, such a modular approach is not well suited to coping with them. And we know that that is true. People catch infectious diseases from other people; and the best way to prevent one patient from becoming infected is often by preventing others around her from becoming infected. That is what "public health" is all about – and public health programmes are, of course, the antithesis of modular one-person-at-a-time approaches to health care.

³² Fried 1978, p. 176.

Not only are there important interactions among patients on the “disease” side of the equation. So too are there important interactions among patients on the “treatment” side. Some of them are positive. Economies of scale can make the unit cost of treating any given patient's condition a decreasing function of the number of other patients afflicted with the same condition.

Some of the more important interactions are negative, however. Insofar as medical resources are scarce relative to the need for them (and they virtually always are), one person's treatment comes to some greater or lesser extent at the cost of some other person's being treated. The modular approach assigns each health care provider responsibility for zealous advocacy of her patients' medical interests – even if at some cost to the greater needs of others who happen not to be her patients. That can easily lead to misallocation of scarce medical resources, unless (contrary to fact, everywhere in the known world) everyone in need of health care had an equally effective advocate whose zeal were strictly proportional to her patients' medical needs.

B. Sector-specific Institutional Responsibilities

An analogous failure due to modularity occurs at the level of intersectoral institutional design. The tasks of government are divided up along

functional lines, and responsibility for different functions is assigned to different departments.

Again, that modular solution would work well were the problems to be solved decomposable in form. But typically they are not. Population health is powerfully affected not only by what happens in Ministries of Health but also by policies laid down by other ministries (of Housing, of Employment, of Environment – not to mention War). Each agency is left maximizing over its own particular remit, which affects performance under others' remits, but in ways that are not its job to take into account.

A classic case of suboptimization ensues, as each maximizes on its own dimension of responsibility without reference to the cost on dimensions that are the responsibilities of others. The Theory of Second-best teaches us that we should generally look at choices in a holistic way, assessing options along all the relevant dimensions at once. By “departmentalizing” policy choice, the modular approach to institutional design does the opposite of that. Each department is assigned some dimension (or small set of dimensions) to be its own particular responsibility, and it is to assess policies within its portfolio under that (those) dimensions.

Were the spillovers across departmental remits either rare or only on high-profile problems, there might be some hope of resolving them through special interdepartmental committees (or, in extremis, at the Cabinet table itself). But there is every reason to think that the spillovers are ubiquitous, and that coordination is poor. Certainly when it comes to regulating health

risks, “coordination... in modern government... has not been pursued in any systematic way.”³³

This has grievous effects on public health in a great many ways. For just one example, reflect upon what is now known about the “social determinants of health.”³⁴ Spending just a little money on good public housing can save a lot of money in terms of keeping people out of hospital. If both those items appeared on the same organization's ledger, that expenditure would be more than balanced by the money it saved. But those items appear on different departments' ledgers. In consequence, the health costs and benefits of policy in other domains are not fully internalized by those other departments with line responsibility for the actions in question. And attempts to overcome that by giving regional authorities responsibility for a wide range of health and community services, and to reallocate resources within that portfolio, have not met with any conspicuous success.³⁵

Another aspect of the institutional carve-up of responsibilities is of course between countries: pollution, transmission of infectious diseases, consequences of trade and labor market policies of one country on the health of citizens in other countries, and so on.³⁶ There are clearly health-related spillovers across national borders. However much cross-border cooperation there may be in mounting coordinated responses to them, it is undeniably the

³³ In the words of Sunstein (1996, p. 1555) – now Administrator of the Office of Information and Regulatory Affairs in the Executive Office of the President of the US who is responsible for doing just that.

³⁴ WHO CSDH 2008.

³⁵ Certainly not in the case of the Prince Edward Island experiment, under the 1993 Health and Community Services Act, anyway (Stoddart et al. 2006).

case that dividing responsibility among different jurisdictions leads to a certain amount of suboptimization as each inevitably prioritizes the health needs of its own citizens over those of others.

C. One Issue at a Time

Herbert Simon was fond of saying that the human mind is essentially a serial processor: it tends to focus attention on one thing at a time.³⁷ When treating patients, physicians would never dream of attending to just “one condition at a time,” ignoring the way a patient's various conditions and treatments might interact. Yet when making public policy on health, we tend to do precisely that.

We are familiar with one-thing-at-a-time thinking on health through the various “wars” serially waged. Just think of the March of Dimes against polio, the War on Cancer, and so on.³⁸ Such one-thing-at-a-time thinking is reflected in public health schemes of “vertical” service delivery in developing countries (such as National Immunization Days) that target specific interventions in ways not fully integrated into the rest of the health system.³⁹ One-thing-at-a-time thinking is seen, too, in the way that even complex omnibus legislation, like the 2010 US health care reform bill, is almost

³⁶ For a sophisticated political-philosophical treatment of the latter issue, see Young (204).

³⁷ Simon 2002b.

invariably discussed in terms of just one lightning-rod feature at a time – rotating serially, in that case, between "death squads," abortion and "the public option."

More generally, time on the legislative calendar is strictly limited. It is conventional wisdom that there is room on the agenda of the UK Parliament for only about 20 major pieces of legislation to be enacted in any given year.⁴⁰ So limits of parliamentary scheduling as well as public attention also ensure that we cannot discuss everything all at once.

But if the policy problems are interrelated (and all are, if only in that they all compete for scarce budgetary resources), then we risk falling afoul of the Theory of Second Best by not doing so. The modular approach of thinking about issues one-at-a-time (or a-few-at-a-time) leads to suboptimization, as we choose policy options that are best on the dimensions we are focusing on without regard to their impact on other dimensions we are not.

This occurs, for example, whenever government departments (of health or anything else) are told to work within some fixed budget. They do the best they can with the resources they have been given. But (the Theory of Second-best teaches us) they might well have done something completely different, had they been allocated either more or less resources.

³⁸ Observing how the War on Terror was displacing the War on Drugs, one character on "The Wire" is lead to ask, "Is our heart not big enough for two wars?"

³⁹ Msuya 2003.

⁴⁰ Mechelen and Rose 1986.

Just as we fall afoul of the Theory of Second Best when taking the budget allocated to health (or anything else) as given, so too do we do so whenever taking the institutional allocation of responsibilities as given. Redrawing the departmental allocation of responsibilities to make the same organization responsible for both health and housing, insofar as those importantly interact, would reduce institutionally-induced suboptimization that comes from different departments focusing on each in isolation. Redrawing the boundaries to make all of sub-Saharan Africa a single decision-making unit would reduce institutional barriers to a consolidated HIV-AIDS strategy for the region.

When discussing the "internalization of externalities," political economists often talk about redrawing the boundaries so as to make a single unit responsible for regulating all activities that affect one another within some policy domain. Making a single decision-making body responsible for an entire watershed – the Rhine, for example – eliminates the risk of some upstream jurisdiction doing something that would make economic sense only if it could count on passing off the costs onto downstream jurisdictions. The case for broadening jurisdictional responsibilities from within the Theory of Second-best reinforces that: doing what is best for the Rhine valley overall is importantly different from doing what is best for the Rhine taken one-stretch-at-a-time. Decomposition does not work well given interactive systems of that sort.

D. Strategies for Overcoming the Errors of Decomposition

When we have erred in trying to decompose problems that are not truly decomposable, run into Second-Best style problems in consequence, there are broadly speaking two standard ways of solving the problem. Both involve taking a more holistic approach to the problem. But the one takes a holistic approach from the bottom-up, while the other takes a holistic approach from the top-down.

Both strategies are nicely illustrated in recent public-service reforms in the UK. For an example of the bottom-up approach to getting a more holistic perspective, consider the innovation whereby GP-fundholders became consolidators with responsibility and resources to provide comprehensively for all the medical needs of their patients. All the interacting factors involved in providing good medical care to a patient were thus placed in a single person's hands. For an example of the top-down approach, consider the "joined-up government" initiative that tried to coordinate to the multiply interacting aspects of public policy at the very center of government.

Neither is a wholly satisfactory approach. The bottom-up approach is (or can be if it works well) good for capturing the interactions among policies impacting any given person; but it is no good for capturing interactions among people. And, as I have already said, it crucially supposes that each person has an equally effective advocate. The top-down approach is (or can be if it works well) good for capturing the interactions among policies within a jurisdiction; but it is no good for capturing interactions among jurisdictions.

Each of those standard approaches has merit, in its place. But something more is needed to solve the widest-scale errors of Second-Best arising from decomposing problems that are not decomposable. For that, we need to implement something more like the "all-encompassing" approach suggested by the "internalizing externalities" story just told. We need to create some authority that is actually responsible for all of the interacting elements taken together. For problems that importantly interact all across the globe, that would mean a world government with responsibility for those problems.

V. Another Partial Solution: Robustness Against Changing Possibilities

So far I have been using the term "second-best" in a generic way to refer to anything that is not first-best. But of course the same sorts of feasibility constraints that force us to fall back on second-best solutions can, if the literal second-best is also infeasible, also drive us to third-, fourth- or fifth-best solutions.

The Theory of Second-best applies all the way down. Just as the second-best solution might be unlike the first-best in every respect, so too the third-best will often be unlike the first- or second-best in every respect. And so on for the fourth-best, et seq.

Now, suppose feasibility constraints change more quickly than policy choices can change. That is a not-uncommon phenomenon in public policy

in general. Path-dependency and historical lock-in are familiar features.⁴¹ That fact, when combined with the Theory of Second-best, has important implications for policy choice.

Suppose we have chosen the third-best alternative, because that was the best available at the time we had to choose. But now suppose feasibility constraints ease, and the second-best alternative suddenly becomes an option. If (per the Theory of Second-best) the second-best alternative is wholly unlike the third-best, and if we are locked into a policy fit only for pursuing that very different third-best alternative, we will be unable to avail ourselves of this new opportunity.

Such reflections give rise to a further prescription for policy-making. Whenever feasibility constraints are likely to change more quickly than policy, the rule ought to be: choose a policy that is robust against changes in feasibility. Thus, in the example just offered, it might be better to opt for a policy to pursue the fourth-best alternative instead of the third-best (which is also genuinely available), if that policy would be more amenable to pursuit of the second- or even first-best alternatives should they become available.

Like all maxims of policy, this one should provide only a pro tanto reason for action. If the fourth-best alternative is very much worse than the third-best, or if it seems very unlikely that the second- or first-best alternatives will become feasible anytime soon, then this maxim might well be overridden. So the robustness strategy, like all the others, is highly

⁴¹ Arthur 1988; 1989. Pierson 2000

sensitive to the empirical facts of the matter, and there is no reason to think it will always prove possible or even desirable. Still, “robustness of policy choice against changing feasibility” is a consideration that ought always be borne in mind in choosing among “sticky” courses of action.

Suppose, for a bioethical example, we are trying to prevent the spread of some infectious disease. The first-best alternative would be for everyone to be vaccinated against the disease. But suppose it is technically infeasible to produce enough vaccine to do that. How then ought we best distribute the limited supply of vaccine?

Imagine that, upon inspection of the interaction patterns among the populations at risk, you notice that they sort themselves into several relatively self-contained communities, with only a few people passing between them. (Think for example of how HIV is spread by long-distance truckers to remote communities in southern Africa.) Given that observation, one strategy would be to try to create herd immunity within at least some of those communities (although only some since – let us suppose – there is not enough vaccine to create herd immunity within all such communities). Another strategy, aimed more at preventing the spread of the disease from one community to others, would be to inoculate everyone who passes between communities (which would also – let us suppose – pretty fully exhaust supplies of the vaccine).

Obviously, once the vaccine has been injected into one person it cannot be extracted and re-injected into another. So once we have implemented one or other of these policies, we are “locked in.” The first condition for applying

my “robustness” rule thus applies. The second (quickly changing possibilities) might as well. Suppose for example new transportation possibilities suddenly emerge, and lots of new people start plying the trade routes between the various communities.

Absent robustness considerations, the ranking of policy options would seem to be that: (1) immunizing everyone at risk is first-best (but ex hypothesi infeasible); (2) immunizing travelers between communities is second-best, because if successful it promises to protect all communities; (3) immunizing enough people to create herd immunity within some communities is third-best, because it leaves other whole communities unprotected.

That is the way things look, without taking robustness considerations into account. Given robustness considerations, however, and the real risk that there might suddenly be so many travelers that that strategy collapses, it might be better to opt for distributing the vaccine so as to give at least some communities herd immunity.

V. Conclusion

The upshot is a moderately discouraging one. The trouble that the Theory of Second-best makes for policy choice is not easily avoided. It issues firm cautions against the strategy of bringing the real world more into line with

the presuppositions of ideal theory, in much the same way and for much the same reasons it does against the more standard strategy of adapting ideal theory to fit the real world. The Theory of Second-best further warns that theorizing in the vicinity of the actual might be highly unstable; that problems might not be decomposable; and that there may be no policy choice that would be robust against all possible (or even all likely) eventualities. In short, if we are looking for guarantees, there seem to be none.

The Theory of Second-best might unsettle any of those strategies. Or it might not. On the weaker and more general form I have been focusing on in this chapter, there is no reason to think it will necessarily always do so. It may turn out that there are some realms in which the second-best really is exactly like the first-best in all respects except one. Whether there are and what they might be is an empirical matter, not resolvable by philosophical analysis alone. The most the philosopher can do is to warn policy-makers to watch out for interactions across evaluative dimensions and to be sensitive to important evaluative dimensions being omitted altogether.

References

- Arthur, W Brian. 1988. Self-reinforcing mechanisms in economics. In Phillip W. Anderson, Kenneth J Arrow & David Pines, eds., The economy as an evolving complex system. Reading, Mass.: Addison-Wesley.
- Arthur, W. Brian. 1989. Competing technologies & lock-in by historical events. Economic Journal, 99: 116-31.
- Barry, Brian. 1965. Political Argument. London: Routledge & Kegan Paul.
- Bentham, Jeremy. 1827. Rationale of judicial evidence. Vols. 6-7 in Works, ed. J. Bowring. Edinburgh, W. Tait, 1843.
- Black, Max. 1962. Models and Metaphors. Ithaca, NY: Cornell University Press.
- Carens, Joseph H. 1996. Realistic and idealistic approaches to the ethics of migration. International Migration Review, 30 (#1: Spring): 156-70.
- Fried, Charles. 1978. Right & Wrong. Cambridge, MA: Harvard University Press.
- Gilbert, Pablo. 2009. The feasibility of basic socioeconomic human rights: a conceptual exploration. Philosophical Quarterly, 59 (#237: Oct): 659-81.
- Goodin, Robert E. 1982. Political Theory & Public Policy. Chicago: University of Chicago Press.
- Goodin, Robert E. 2009. Demandingness as a virtue. Journal of Ethics, 13 (#1: Mar), 1-13.

- Healey, Kieran. 2000. Embedded altruism: blood collection regimes in the European Union. American Journal of Sociology 105 (#5: Mar): 1633-57.
- Healey, Kieran. 2006. Last Best Gifts: Altruism and the Market for Human Blood and Organs. Chicago: University of Chicago Press.
- Lawford-Smith, Holly. 2010. Debate: ideal theory – a reply to Valentini. Journal of Political Philosophy, forthcoming.
- Lipsey, Richard G. and Kelvin J. Lancaster. 1956. The general theory of second best. Review of Economic Studies 24: 11-33.
- Margalit, Avishai. 2010. On Compromise and Rotten Compromises. Princeton, N.J.: Princeton University Press.
- Mechelen, Denis van and Richard Rose. 1986. Law as a resource of public policy. Parliamentary Affairs, 36 (3): 297-314
- Msuya, Joyce. 2003. Horizontal and vertical delivery of health services: what are the trade offs? No. Background Study 26942 for World Development Report 2004/5. Washington, D.C.: World Bank.
Available at < http://www-wds.worldbank.org/external/default/WDSContentServer/IW3P/IB/2003/10/15/000160016_20031015125129/additional/310436360_200502761000211.pdf>.
- Mulgan, Timothy P. 2001. The Demands of Consequentialism. Oxford: Oxford University Press.
- Murphy, Liam B. 2000. Moral Demands in Non-ideal Theory. Oxford: Oxford University Press.

- Nelson, Aaron L.; Joshua T. Cohen; Dan Greenberg; and David M. Kent.
2009. Much cheaper, almost as good: decrementally cost-effective medical intervention. Annals of Internal Medicine, 151 (9: 3 Nov): 662-7
- Okun, Arthur. 1975. Equality and Efficiency: The Big Tradeoff. Washington, D.C.: Brookings Institution.
- Pierson, Paul. 2000. Increasing returns, path dependence and the study of politics. American Political Science Review, 94 (#2: June): 251-68.
- Rawls, John. 1971. A Theory of Justice. Cambridge, MA: Harvard University Press.
- Rawls, John. 1993. Political Liberalism. New York: Columbia University Press.
- Rawls, John. 1999. The Law of Peoples. Cambridge, Mass.: Harvard University Press.
- Rawls, John. 2001. Justice as Fairness: A Restatement. Cambridge, Mass.: Harvard University Press.
- Rousseau, Jean-Jacques. 1762. The social contract. Pp. 164-278 in The Social Contract and Discourses, trans. G.D.H. Cole. New ed. London: Everyman/Dent, 1973.
- Simmons, A. John. 2010. Ideal and nonideal theory. Philosophy & Public Affairs, 38 (1): 5-36.
- Simon, Herbert A. 1969. The Sciences of the Artificial. Cambridge, MA: MIT Press.

- Simon, Herbert A. 2002a. Near decomposability and the speed of evolution. Industrial and Corporate Change, 11 (#3: June): 587-99.
- Simon, Herbert A. 2002b. Organizing and coordinating talk and silence in organizations. Industrial and Corporate Change, 11 (#3: June): 611-8.
- Stoddart, Greg L.; John D. Eyles; John N. Lavis; and Paul C. Chaulk. 2006. Reallocating resources across public sectors to improve population health. Pp. 327-47 in Healthier Societies: From Analysis to Action, ed. Jody Heymann, Clude Hertzman, Morris L. Barer and Robert G. Evans. Oxford: Oxford University Press.
- Sunstein, Cass R. 1996. Health-health tradeoffs. University of Chicago Law Review, 63 (#4: Fall): 1533-72.
- Valentini, Laura. 2009. On the apparent paradox of ideal theory. Journal of Political Philosophy, 17 (#3: Sept): 332-55.
- Weber, Max. 1919/2004. Politics as a vocation. Pp. 32-94 in Weber, The Vocation Lectures, ed. David Owen and Tracy B. Strong, trans. Rodney Livingstone. Indianapolis, IN: Hackett.
- WHO CSDH. 2008. Closing the Gap in a Generation: Health Equity through Action on the Social Determinants of Health. Final Report of the Commission on Social Determinants of Health. Geneva: World Health Organization.
- Young, Iris Marion. 2004. Responsibility and global labor justice. Journal of Political Philosophy, 12 (#4: Dec): 365-88.
- Ypi, Lea. 2010. On the confusion between ideal and non-ideal in recent debates on global justice. Political Studies, forthcoming.

