May a Government Mandate more Comprehensive Health Insurance than Citizens Want for Themselves?[1]

Alex Voorhoeve, LSE and NIH

Abstract

I critically examine a common liberal egalitarian view about the justification for, and proper content of, mandatory health insurance. This view holds that a mandate is justified because it is the best way to ensure that those in poor health gain health insurance on equitable terms. It also holds that a government should mandate what a representative prudent individual would purchase for themselves if they were placed in fair conditions of choice.

I argue that this common justification for a mandate is incomplete. A further reason for mandated insurance is that it helps secure social egalitarian public goods that would be underprovided if insurance were optional. I also argue that rather than mandating what a representative individual would choose for themselves, we should design the mandatory package by appealing to a pluralistic egalitarian view, which cares about improving people's well-being, reducing unfair inequality, and maintaining egalitarian social relations.

"What do you say to Mark and Lucinda in my district who had a plan, they liked it, it was affordable, but it is being terminated [because it is not sufficiently comprehensive to meet the requirements of the Affordable Care Act]? It was what they wanted and I will remind you: some people like to drive a Ford, not a Ferrari. Some people like to drink out of a Red Solo cup, not out of crystal. You are taking away their choice."

<div align="right">

Representative Marsha Blackburn (Republican) to U.S. Secretary of Health Kathleen

Sebelius, Hearing on the Affordable Care Act, 30 October 2013

</div>

Introduction

In most developed and many developing countries, governments require citizens to be insured for a minimum set of health care interventions. They typically do so through some combination of the following: (i) the provision of tax-financed care (as, for example, Medicare does for the elderly in the U.S.); and (ii) requiring the purchase of at least a basic package of insurance from a social fund or a private provider (as the U.S. Affordable Care Act does for the non-poor and non-elderly) (Carrin and James 2005; Cotlear *et al*. 2015). For simplicity, in what follows, I shall refer to both tax-financed care and required purchase as "mandated insurance."

In this paper, I shall address two questions about such a mandate. First: what, if anything, justifies the limitation it places on the freedom to spend one's fair share of resources as one sees fit? Consider, for example, the typical citizens invoked by Representative Blackburn. Suppose they have average incomes and modest current health risks, are prudent and adequately informed. Why shouldn't they be able to spend the money they are required by the Affordable Care Act to spend on health insurance in whatever way they prefer, for

example by purchasing cheaper, less comprehensive insurance and using the remaining funds for other purposes? The second question is this: if mandatory insurance is justified, how should governments determine the content of the mandated minimum package?

Various answers to these questions have been proposed, in part because mandated coverage serves a variety of ends (see Daniels 2013 for a review). One such end is protection against the consequences of imprudent risk-taking (Dworkin 2000, p. 492n6; Bou-Habib 2006; Moncrieff 2013). Another is to fairly and efficiently cover the costs of the care that we are obligated to provide on humanitarian grounds, such as basic care for those in urgent need who could not afford to pay for such care out of pocket (Buchanan 1984; Bou-Habib 2006; Segall 2010, chap. 5; Menzel 2012; Rulli and Wendler 2016). In this paper, I shall not focus on these reasons for requiring coverage, important though they may be. Instead, I shall critically analyze the following well-known pair of liberal egalitarian answers for the case of prudent, adequately informed adults (Dworkin 2000, chap. 8, 2012a,b; Braun 2012; Menzel, 2012):[2]

> **(i). To keep premiums affordable for those at unfairly high risk, those at low risk must be forced to insure themselves on relatively equal terms.**
>
> Fairness towards those who, due to bad brute luck, have comparatively large anticipated health care needs requires that they be permitted to insure themselves on similar terms as those with smaller anticipated needs. But insurance on such equal terms will be most attractive to those with the greatest health care needs and unattractive to those with the smallest needs. If such insurance were voluntary, therefore, a disproportionate share of enrollees would be those with relatively high health care costs. Coverage for this group would be expensive, and there would be

3

little of the desired sharing of risks and costs between the lucky and unlucky. This outcome can be most efficiently prevented by forcing those with low expected health care costs to insure themselves on terms that are similar to the terms for those with high expected costs.

(ii.) The mandated minimum package of services should be determined by the Representative Prudent Individual Test (or RPIT, for short).

A government should mandate the insurance package that a typical prudent and adequately informed individual would wish to purchase for themselves if they were placed in fair conditions of choice.

In this paper, I will argue that answer (i) is incomplete. A further reason to mandate insurance is that this contributes to important social egalitarian public goods. I will also argue that answer (ii) is insufficiently egalitarian.

My argument proceeds as follows. In Section 1, I offer a fuller explanation of the common liberal egalitarian justifications for a mandate and the RPIT. In Section 2, I consider circumstances in which individuals face differential risks. I argue that under such circumstances, mandating the insurance that a representative prudent individual would purchase wrongly treats *inter*personal trade-offs as if they were *intra*personal trade-offs and thereby ignores unfair inequalities. In Section 3, I consider circumstances in which all individuals are at equal risk and have equal purchasing power and explain why, under such circumstances, the common liberal egalitarian account holds that there is no reason for a mandate. In Section 4, I argue that this account fails to consider the social egalitarian public goods to which a well-designed mandatory health insurance package contributes. These public goods are underprovided when individuals are free to make their insurance decisions

independently, and this, I argue, gives us a reason to mandate insurance even when people are equally situated. I conclude that governments should design a mandatory package by appealing to a pluralistic egalitarian view, which cares about improving people's well-being, reducing unfair inequality, and maintaining egalitarian social relationships.

1. The argument for mandating the package desired by a representative prudent individual.

Leading thinkers have argued that a mandated minimum of insurance is justified as the best solution to two problems that would occur in a free market for health insurance (Gibbard 1984; Dworkin 1994, 2000, chap. 8, 2012a,b; Braun 2012; Menzel 2012; WHO 2014).

The first is that people with low expected expenditures on needed health care in a given period would be charged relatively little for insurance in that period, while those who would have high expected expenditures would be charged much more. Assume, as is plausible, both that those with higher expected expenditures tend to be those with worse health and that a large proportion of differential health is not the result of free, adequately informed choice under fair circumstances but is rather due to bad brute luck. It follows that a free market in health insurance would add financial burdens to the unfair disadvantage of poor health. Such compounding of burdens would be especially problematic because those with worse health are already more likely to have lower incomes and lower social status (Deaton 2013; Sreenivasan 2014).

One standard way of avoiding such unfair premium differentiation is to require that insurers take anyone who applies at a premium that is relatively independent of the applicant's personal risk profile. (This is known as "community rating.") However, such community

rating contributes to a second problem. Suppose that firms were to set the uniform premium at a price that would cover the average person's health expenditure. At this premium, insurance would be least attractive to people who believe themselves to have lower-than-average expected expenditures on covered health services and most attractive to people with higher-than-average expected expenditures. Since many people are aware of their (approximate) risk profile, those who would enroll would tend to have disproportionately high health expenditures, a phenomenon known as "adverse selection" (Arrow 1963). To cover the costs of this disproportionately sick, enrolled population, firms would therefore have to raise the premium above the initial rate. But doing so would, at the margin, tend to cause those with the lowest anticipated health costs to drop out of the pool, since they would then regard insurance as too expensive relative to their personal risk profile. This would further worsen the risk profile of the enrolled population, meaning that the premium would have to be raised again, leading to a so-called "death spiral" of an escalating premium and an insurance pool consisting of people with ever-greater health risks. Without intervention, this adverse selection mechanism may severely impair the functioning of the insurance market (Akerlof 1970; for a real-world example, see Barro 2010).

In theory, these problems might be addressed by pairing community rating with extensive subsidies for voluntary enrolment. However, the experience of countries at all levels of development suggests that this is not a promising route (Carrin and James 2005; Lagomarsino *et al*. 2012; Cotlear *et al.* 2015). By contrast, together, community rating and mandatory enrolment achieve both fair cost-sharing between those at low and high health risk and a solution to the adverse selection problem.

Requiring enrolment of those who would rather not now enroll because they are currently at low risk is not merely a case of forced transfers to others who are less well off. Presumably, many people who are currently at low risk want health insurance available to them at reasonable cost when their risks become high (such as when they develop a need for expensive care). Due to the aforementioned death spiral, insurance might not then be available to them at affordable cost if enrolment were voluntary. Many may therefore prefer, on their own behalf, that everyone is required to be insured when they are at low risk in order to guarantee that there is a well-functioning insurance market with reasonably priced coverage when their expected health care costs become high. Mandatory payments may be justified to these people as an efficient way to get something they would be unable to get if everyone were free to make insurance decisions on their own (Braun 2012; Menzel 2012).

Naturally, to avoid unfairly burdening the poor, contributions (through taxes and/or individual premiums net of subsidies) should be income-dependent. A mandate of this kind redistributes resources from rich to poor and from the healthy to the infirm. Like many solidaristic arrangements, it has a mixed justification. In part, it appeals to an egalitarian principle of justice that cares about reducing unfair inequality. In part, it appeals to each person's self-interest in establishing an institution that serves this interest better than uncoordinated individual choices would.

I now turn to the question how one ought to draw the boundaries of the package of mandated care. To many, the following line of reasoning has seemed compelling (Gibbard 1984; Dworkin 1994, sec. 2; 2000, ch. 8; Menzel 2012, p. 596n28; Kurtulmus 2012).

Suppose that each member of our society were prudentially rational, had a fair share of resources and the average health risks of people in our society. Suppose further that they had to choose a lifetime health insurance policy with full pertinent information about the prospective benefits and costs to them of each feasible package of health services. In such circumstances, there would be no need to require community rating, since everyone's health risks (and their beliefs about these risks) would be identical at the moment the insurance contract was signed. For the same reason, there would be no need to prevent adverse selection. Moreover, there would be no other reasons of justice to interfere with individuals' choices regarding health insurance—the resulting pattern of insurance would be fully just. We should use this idealization of our society as a reference point for mandatory insurance in our actual, unjust circumstances. In the idealized circumstances, individuals with different preferences would purchase different insurance plans. We cannot replicate this diversity of individually-tailored insurance plans in our society for many reasons, including the adverse selection problem outlined. But we should approximate it by determining the content of the required insurance package as follows.

The Representative Prudent Individual Test (RPIT): We should mandate the lifetime package of insurance that a representative prudentially individual with a fair share of resources would want for themselves if: (a) their self-regarding preferences and risk attitudes were those of a typical person in our society (once corrected for irrationalities); (b) they believed that they faced the average lifetime health risks in our society; (c) they knew the potential costs and benefits to them of the package; and (d) their risks, costs and benefits were calculated under the assumption that the package selected would become the required minimum for all.[3]

In the remainder of this paper, I shall raise objections to this view and propose revisions to it.

2. Objections to the RPIT

I shall begin by arguing that the RPIT needs to be discarded because it is insufficiently egalitarian. The following Differential Prospects Case highlights its problems. Suppose that a three-person society of UNHEALTHY, HEALTHY, and AVY is about to implement a new mandatory package. UNHEALTHY has recently been diagnosed with a serious health condition and her lifetime need for health care is certain to be great. HEALTHY, by contrast, is certain to possess excellent lifetime health and have little need for health care.[4] AVY has average prospects: she is in good health now, but will end up in either the same situation as UNHEALTHY or the same situation as HEALTHY, with each possibility being equally likely. The three feasible policies are: Small (a low-cost package with limited coverage), Medium (a moderately costly package of middling coverage), and Large (a more costly, but also more comprehensive package). So that we have a concrete measure of prudential interests in terms of which we can assess these packages, I shall assume that prudential interest is identical to well-being and is measured as follows, in line with orthodox decision theory.[5] A first alternative has higher expected well-being for a person than a second alternative just in case the first would be strictly preferred for this person's sake after rational, calm deliberation with all pertinent information while considering their self-interest only. Two alternatives yield equal expected well-being just in case such deliberation would yield indifference between them.[6] The assumed impacts of these three insurance packages on each person's well-being are represented in Table 1.

Table 1. Lifetime well-being for three insurance packages in the Differential Prospects Case.

Individual Policy	UNHEALTHY	AVY	HEALTHY
Small	35	59.5	84
Medium	40	60	80
Large	44	59.5	75

Note: UNHEALTHY's and HEALTHY's well-being is certain under each policy; the well-being listed for AVY is her expected well-being.

In this case, by construction, AVY's prospects are equally good under Small and Large. The RPIT is therefore indifferent between them. Moreover, her prospects are most advanced by the choice of Medium. Since AVY is the representative individual, the RPIT therefore selects this package.

The central objection to this test is that it treats what is a trade-off between the interests of two separate people, UNHEALTHY and HEALTHY, in precisely the same manner as it treats a trade-off between two potential futures of a single individual, AVY. In doing so, it fails in two ways to recognize reasons for favoring the worse off that arise in interpersonal trade-offs (Gauthier 1963, pp. 121-7, Nagel 1970, pp. 132-42; Rawls 1999, secs. 5 and 39).

First, the RPIT is insensitive to the fact that some will be better off than others. When one evaluates a policy solely on the basis of AVY's expected well-being, one pays attention to

how she will fare under that policy in each of her potential futures, but one does not consider it as in itself bad that in one potential future she would be better off than she would be in another potential future. Small and Large are equally good prospects for Avy because they yield the same expected well-being, even though the gap between her worst and best potential outcome under Small is greater than the gap between her worst and best potential outcome under Large. The RPIT takes the same attitude to the fates of co-existing, separate people as it takes to the equally likely potential futures of a single person. It is therefore indifferent to the inequality between two coexisting, separate people (UNHEALTHY and HEALTHY), just as it is indifferent to the inequality between two mutually exclusive possible futures of the same person (AVY). For the RPIT, Small and Large are equally good policies, even though the former contains more interpersonal inequality than the latter. This lack of concern for how some fare compared to others is not merely substantively mistaken; it is also inconsistent with the aversion to unfair inequality that motivates the common liberal egalitarian account.

Second, the RPIT fails to recognize that when different people's interests conflict, the strength of a person's claim to a benefit of a given size depends in part on how badly off this person is compared to others.[7] By way of illustration, consider again Small and Large. As just noted, the RPIT holds that these two policies are equally good, so that, if one faced a choice between these two alone, it would be permissible to choose either one. UNHEALTHY could challenge this judgment with the following rhetorical question (Otsuka and Voorhoeve 2009, pp. 183-4):

> "How could you choose to forgo a very substantial benefit to a worse off person to instead provide an equally large benefit to someone else who will, in any case, be much better off?"

Next, consider the choice between Small, Medium, and Large. The following complaint on UNHEALTHY's behalf against a choice of Medium rather than Large would be almost as forceful, and, I submit, equally unanswerable:

> "How could you choose to forgo a very substantial benefit to a worse off person to instead provide only: (i) a somewhat larger benefit to a different, much better off person and (ii) a comparatively small improvement in the expectations of another person, who already has far superior prospects?"

No such questions arise when we compare these policies solely in terms of their impact on AVY's prospects—one could not complain, on behalf of the unlucky possible future of AVY that by choosing Medium (rather than Large) one would be failing to benefit them in order to provide a somewhat larger benefit to a different person (Otsuka and Voorhoeve 2009, p. 184).

These problems are avoided if we abandon the RPIT and determine the content of the mandatory package with reference to a pluralist egalitarian view, which cares about reducing unfair inequality in both prospects and outcomes and about improving total well-being.[8] On such a view, Large will obviously be regarded as better than Small, since Large contains less inequality than, but the same total expected well-being as, Small. Properly specified, this pluralist egalitarian view will also rank Large over Medium, because the former very substantially reduces inequality at only a minor cost in terms of total expected well-being. Moreover, because this view is averse to inequality, it recognizes that

12

UNHEALTHY's claim to Large is strengthened by the fact that she is less well off than others and that HEALTHY's claim against Large is weakened by the fact that she is better off than others. Finally, using this form of pluralist egalitarianism to design the mandatory package also makes for a consistent rationale, because it means that the very same values that motivate the mandate—reducing unfair disadvantage and improving the general welfare—also determine what is mandated.

It is worth emphasizing that while, in the context of this simple example, such pluralist egalitarianism will recommend Large because it improves the situation of the worst off at an acceptable cost, it will not turn health care into a "bottomless pit." One reason is that, plausibly specified, this pluralist view will rule that a small improvement in the situation of the worst off is not worth a very large diminution of others' prospects. To illustrate, suppose that in our simple case, the government could also mandate an Extra Large insurance package, which would improve UNHEALTHY's lot (and, if she were unlucky, AVY's as well) to a well-being level of 45, but at the cost of leaving HEALTHY (and, if she were lucky, AVY) with a well-being level of 65. A reasonable pluralist egalitarian view will, I submit, prefer Large to Extra Large, because the reduction in total expected well-being that the latter would involve would be too great compared to the improvement in the prospects of the least well off.

Moreover, the real world is more complex than our example, in which more comprehensive coverage always leads to greater transfers from the healthier and wealthier to the sicker and poorer. This example is useful, because it demonstrates the RPIT's failings and because such transfers are an important part of well-designed, real-world mandatory insurance packages. And insofar as, in the real world, an expansion of coverage leads to more such transfers, the proposed egalitarian criterion will indeed require more comprehensive

coverage than the RPIT. But in reality, there may also be interventions that would modestly improve average well-being but also disproportionately benefit the better off and impose a substantial burden on the worse off, either in financial terms or in terms of other improvements to their lives that could be made with the resources in question. While the RPIT would recommend including these interventions, the pluralist egalitarian view would not. The pluralist egalitarian package will therefore likely include both more coverage for some conditions (e.g. those that primarily afflict the worst off) and less coverage for others (e.g. those that primarily affect the better off).

More generally, since it requires that finite resources be spent in the manner that best advances the aims of improving the general welfare and reducing inequality, the view proposed here will motivate cost-control measures in health care, such as those that reduce the frequency of very expensive interventions with low expected benefit (Menzel 2012, pp. 591-4). It will also direct us not to spend money on health when we could do more to improve people's lives and reduce inequality by spending it on education, environmental quality, personal consumption, and so on.

In sum, the proposed pluralist brute luck egalitarian view can avoid the objections I have raised against the RPIT while taking account of the need to limit expenditure on health. As we shall see in the next two sections, however, the outlined view is incomplete, since it needs to be supplemented by a principle that captures distinctive social egalitarian concerns.

3. The case of identical risks

In this section and the following, I shall argue that the grounds for a mandate extend beyond the need to solve the problems created by unfair differential risks and adverse selection. I shall therefore focus on a scenario in which these problems are absent.

Consider the following Identical Prospects Case. Suppose a large society consists only of adult Aᴠʏs, each of whom, as above, faces a 50-50 gamble between a serious illness and a long, healthy life. Also suppose that it is known that precisely half of them will end up in either position. As in our earlier discussion, everyone is prudent and informed of their risks. The government can either leave everyone free to choose between the aforementioned three packages, or instead require everyone to purchase Large. Table 2 represents, for each insurance package, the prospects of each person and the distribution of final well-being in the population that would result if everyone had this package of insurance.

Table 2. Prospects and outcomes for three packages in the Identical Prospects Case

Package	Each person's expected well-being	Distribution of final well-being if everyone had the package
Small	59.5	Half at 35 Half at 84
Medium	60	Half at 40 Half at 80
Large	59.5	Half at 44 Half at 75

Given our assumptions, if each person had a free choice between these insurance packages and acted independently on their own behalf, each person would purchase Medium, because this is what maximizes their expected well-being. This will lead to substantially more outcome inequality than would arise if everyone were instead required to purchase Large. This inequality notwithstanding, in a case such as this, the account I have been criticizing holds that we have no reason to mandate insurance. In Ronald Dworkin's words:

> "whatever this … community actually spends on health care [as the result of people's free choices from a position of equality] is the morally appropriate amount for it to spend: it could not be criticized, on grounds of justice, for spending either too much or too little" (2000, pp. 312-3).

There appear to be several good reasons for Dworkin's conclusion. First, leaving people the choice between these three packages respects their freedom to use an equal share of resources as they see fit. By contrast, mandating Large in order to lessen outcome inequality would appear to infringe this freedom.

Second, given our assumptions, each will use this freedom to make what is, given the risks they face, the prudentially best decision. Concern for the well-being of each person, considered in isolation, should therefore motivate us to endorse each person's choice of Medium; it also counts against mandating Large.

Third, it appears that brute luck egalitarianism does not give us grounds to object to the difference in outcome inequality between Medium and Large. Brute luck egalitarianism objects to unchosen inequalities because they are unfair. It also holds that inequalities that are due to competent, free, and informed choice against an equal background need not be unfair (see, e.g., Cohen 1989; Arneson 1989; Temkin 2001). In the Identical Prospects Case, the increase in outcome inequality between Medium and Large would be entirely due to such choices. This increase in inequality would, of course, be due to luck, but this would not be brute luck, but rather what Dworkin (2000, p. 73) defines as "option luck": "a matter of how deliberate and calculated gambles turn out—whether someone gains or loses through accepting an isolated risk he or she should have anticipated and might have declined."[9] In other words, both Medium and Large would contain the same amount of objectionable brute luck inequality.

Still, it does not follow that a brute luck egalitarian must accept the additional outcome inequality that results from people's choices in this case. By itself, an aversion to brute luck inequality does not tell us what kinds of choice-based inequalities one should permit among

people who face equal opportunities. Brute luck egalitarians must therefore rely on distinct principles to decide which choice-based inequalities to countenance (Olsaretti 2009; Stemplowska 2009; Segall 2010, chap. 4).

One such principle, advanced by Dworkin, is that, in the absence of market failures, the government should allow equally-situated, informed, and rational people to freely contract and let them bear the consequences of these contracts (Dworkin 2000, ch. 2). On the assumption that, with adverse selection and its attendant problems ruled out, there are no market failures in the Identical Prospects Case, it follows that people should be left free to choose and that the inequality generated by everyone's choice of Medium should be accepted.

Another such principle, advanced by Peter Vallentyne, focuses on the value of people's opportunity sets. Vallentyne (2002, pp. 551-2) proposes that when, under conditions of equal opportunity, we must decide whether individuals should face a set of choices that will generate inequality, we should ask whether it would be prudentially valuable for each person to face these choices. Since we are assuming that individuals choose what has higher prudential value for them, this principle also favors leaving everyone a free choice in the Identical Prospects Case.

We have seen that respect for freedom and concern for people's well-being appear to favor leaving people a choice in this case, while the resultant increase in inequality (compared to mandating Large) can be put down to acceptable differential option luck. Despite its appeal, I shall argue in the next section that this line of reasoning is mistaken.

4. Social egalitarian public goods

As an initial step towards this critique, we should note that an individual in the Identical Prospects Case might, while preferring Medium for themselves on prudential grounds, also object to the inequality that results from *everyone* choosing Medium. That is, they might, when considering both their personal prospects and the overall degree of inequality in society, prefer a society in which everyone has Large over a society in which everyone has Medium.[10]

Here, I shall explore the reasons for such a preference that derive from a social egalitarian ideal. This ideal is fully realized when all adult members of society have the ability to live as free citizens[11] who can participate in social life and contribute to public decision-making on equal terms. More specifically, this view involves the attainment of four interrelated goods, each of which is partly constitutive of a society of equals (Anderson 1999; Rawls 2001, sec. 39; O'Neill 2008; Satz 2010; Fourie *et al*. 2015).

The first is the absence of domination and exploitation in private life. Health insurance contributes to this good because it improves people's ability to access needed care without falling into destitution. It thereby prevents the forms of subordination and unfair advantage-taking that occur when the less well-off become highly dependent for needed care or financial assistance on more powerful others.

The second is the absence of domination in political life. To realize this good, society must prevent situations in which those who are less well off do not have their interests effectively represented. Health insurance helps do so when, through transfers from the healthy and wealthy to the unhealthy and less well off, it blunts two of the forces that lead people to

become socially and politically marginalized, to wit, ill health and financial disadvantage (Mattila *et al.* 2013; Smets and van Ham 2013).

The third category encompasses what Rawls (1999) called the social bases of self-respect. Among these is the ability to participate in public life with dignity—to be able to join in work, play, and general social interaction without shame or a sense of inferiority based on characteristics irrelevant to the status of citizen, including low income or ill health. The care to which health insurance gives access can ensure that fewer people have conditions that generate a debilitating sense of shame or inferiority. Moreover, insofar as it redistributes from the well-to-do and healthy to the poor and ill, it limits the type of inequalities that give rise to norms of appearance and functioning that the less-well-off struggle to meet. The care to which it gives access can also bolster other sources of self-worth, including the ability to carry out everyday tasks, as well as to support oneself and one's family.

The final social egalitarian good consists in the attitudes required for social cooperation, which include a willingness to work with others to mutual benefit, to enter into reasoned discussion with them about fair social arrangements, and to abide by fair arrangements if others do so too. Large inequalities in health and income threaten these attitudes in various ways. For one, they can generate a sense of inferiority among the less fortunate and arrogance among the fortunate (Tawney 1964, pp. 37-8). The less well off may also become resentful of others' advantages, leading them to prefer a situation in which the lives of the better off are diminished, even if this does not improve their own lot. In response, the better off may be inclined to jealously guard their relative advantage at a cost to themselves and others (Rawls 1999, pp. 467-8). One further possible effect is described by Adam Smith. The better off tend to avoid interacting with those who are very badly off in order to

prevent the painful sympathetic feelings that would be a natural consequence of being in their company. For their part, an uncomfortable awareness of others' desire to turn away from them causes the badly off to withdraw from society (Smith 1982, I.iii.1; see also Chase and Bantebya-Kyomuhendo 2015). Insofar as it improves the fate of the least fortunate and lessens the gap between them and the rest of society, more extensive health insurance plays a part in preventing the spread of such problematic attitudes.

The ways that health insurance contributes to these goods can provide grounds for a citizen to prefer a society in which everyone has Large over one in which everyone has Medium. But they may also make one wonder whether it is coherent to assume that Medium is prudentially superior. After all, each individual's self-interested prospects should take into account the impact on their well-being of these social effects of an insurance package. These effects might indeed be substantial. Being dominated, politically marginalized, having one's self-respect threatened, feeling envious, or withdrawing from social interaction out of shame all depress one's well-being. A prudent chooser will, therefore, consider the ways in which purchasing Medium (rather than Large) raises the risk that these evils will befall them. Moreover, insofar as we consider the impact of everyone choosing Medium, we should note that an increase in social inequality could depress well-being by lessening social cohesion. Nonetheless, social egalitarian goods (or the lack thereof) are but one component of well-being. It is therefore perfectly possible that from the perspective of each person's self-interested prospects, these potential downsides of Medium are outweighed by its possible upside, namely, being much better off if one were to be healthy. More generally, the importance of social equality cannot be reduced to its prudential value. Independently of one's own well-being, one has reason to pursue a society in which people can relate to one another as equal citizens because such a society involves relationships and attitudes which

21

are in themselves good, and because it limits the evils of domination, indignity, lack of self-respect, and unsociable attitudes. The assumption that Large is prudentially costly but superior in social egalitarian terms allows us to keep the possible trade-off between prudence and social equality clearly in view.

Let us now consider the significance of these social goods and evils for the question of mandating health insurance. Suppose that, considering both their prospects and the degree of social equality, *all* citizens prefer a society in which everyone has Large over one in which everyone has Medium. Wouldn't it follow that if given the independent choice between Small, Medium and Large, each would voluntarily purchase Large?

No. For these social egalitarian goods possess two characteristics of public goods.[12] The first is non-rivalry: one person's enjoyment of an egalitarian environment does not impair others' access to it. The second is substantial non-excludability: if, for example, everyone else were to purchase Large and thereby help secure a society of equals, then whether or not one personally purchased such insurance, one would partake in many of the relevant social goods (e.g. one would live in a society in which domination and exploitation are less common, and in which people are more likely to have the attitudes required for social cooperation).

Together, non-rivalry and non-excludability create a familiar problem in generating the good of an egalitarian environment through independent action. To illustrate this point, consider the simplified representation in Table 3 of the situation of a citizen who must choose between insurance packages independently. (It is simplified by leaving out Small, which is inferior on both prudential and egalitarian grounds, and by considering only two patterns of choice by all other citizens: everyone else chooses Medium and everyone else chooses

Large.) Each cell describes the two things this citizen cares about: their expected well-being and the degree of social equality. By assumption, this citizen prefers the South-East cell (everyone has Large) to the North-West cell (everyone has Medium). But they cannot choose between these two cells, because they control only their own choice of Medium or Large. To choose rationally, this citizen will therefore consider their best reply to what others might do. Now, if all others were to choose Medium, then it is consistent with our assumptions about this citizen's values that they would also choose Medium. For if they were to choose Large, they would only fractionally improve social equality, but they would suffer a significant loss in expected well-being, not merely because of the cost of the insurance package, but also because they would, in any case, suffer the negative effects of the inegalitarian environment created by others' choices. (The latter explains why their expected well-being if *they alone* choose Large is at best equal to their expected well-being if *everyone* were to choose Large.) If, instead, all others were to choose Large, then it is again consistent with our assumptions that this citizen would choose Medium. For their choice of Medium would only fractionally lower the degree of social equality. But it would significantly improve their personal prospects. For they would not have to pay the cost of Large and would benefit from the positive effects on their well-being of living in the more egalitarian society created by others' choice of Large. (The latter are why their expected well-being if *they alone* choose Medium is at least as great as their expected well-being if *everyone* were to choose Medium.)

One might, at this point, object that a citizen who adequately values social equality would be willing to do their bit by purchasing Large if they were confident that enough others would do so too. But we should note that even such a conditional cooperator has every reason to choose Medium if they believe that many others are likely to do so as well.

In sum, even citizens who care enough about social equality to give up some prospective well-being for its sake may find themselves facing a collective action problem, because each may rationally purchase Medium when they face the decision on their own, even though they would prefer a situation in which everyone has Large to one in which everyone has Medium.[13]

Table 3. Relevant outcomes for a citizen choosing between Medium and Large for themselves alone.

All other citizens One citizen	Medium	Large
Medium	Expected well-being = 60 Low degree of social equality	Expected well-being ≥ 60 Degree of social equality is fractionally below 'high'
Large	Expected well-being ≤ 59.5 Degree of social equality is fractionally above 'low'	Expected well-being = 59.5 High degree of social equality

As with collective action problems generally, one solution is to restrict choice, and mandate Large. Especially in a numerous society, such a mandate may be the most efficient solution and may therefore be favored by each member of the population over less restrictive alternatives, such as legally permitting choice but establishing strong social norms in favor of purchasing Large. It is noteworthy that when the mandate is endorsed by all, we do not

have a conflict of values, with on one side, the reasons for permitting choice enumerated in Section 3 and, on the other, social egalitarian concerns. For each is then, by assumption, happy to suffer a small loss in their expected well-being and a limitation of their freedom for the sake of acting collectively to improve the degree of social equality. Moreover, under these circumstances, the two arguments presented in Section 3 for the acceptability of the inequality generated by free choice are not compelling. Recall that the first of these arguments, derived from Dworkin (2000, chaps. 2 and 8), relied on the principle that, in the absence of market failures, equally-situated individuals should be left free to contract, along with the assumption that there are indeed no such market failures in the Identical Prospects Case. We have seen, however, that the latter assumption is false, because the market undersupplies social egalitarian public goods. Dworkin's principle therefore does not imply that we should countenance the inequalities that would result from unconstrained choice.[14] The second argument relied on the principle that inequalities are acceptable when people face equal and prudentially valuable opportunities (Vallentyne 2002). But this principle is questionable in this context, because each citizen will deny that a choice between Small, Medium and Large is, all things considered, the most valuable one they could face. Considering both prudential and social egalitarian value, they prefer not to face this choice (so long as this choice is removed from everyone). And there does not seem to be a strong case for permitting inequalities that result from unwanted choices.

We can conclude that, when all agree to a mandate, there is decisive reason for mandating Large. However, the case for a mandate is less straightforward when some (but not all) citizens would object to it, because they judge that the gain in social equality is not worth the loss of liberty and the cost to their expected well-being. These dissenters need not be purely self-interested; they might merely make a different judgment about the balance of

values involved. In these circumstances, we face a conflict between, on one side, the liberty and prudential interests of those who oppose a mandate and, on the other side, egalitarian public goods that cannot be generated through unconstrained individual choice.

Here, I shall not add to extant discussions about how democratic societies can, in the face of such conflicts of value and opinion, determine a right level of public goods provision (see Claassen 2013 for a review). Instead, I merely note that the arguments made in this section explain why those who favor more comprehensive insurance because of its contribution to social equality cannot be expected by dissenters to simply advance this end through voluntary, private action. That an egalitarian environment is a public good gives us a reason, albeit a defeasible one, to secure it through a mandate, when this is the most efficient way of overcoming the collective action problems that arise in producing and protecting such an environment.

The ways in which health insurance can further social equality are relevant not merely to the justification of a mandate, but also to its content, including in circumstances in which citizens have differential prospects. For these contributions give us reason to judge this package by three criteria: improving total expected well-being, reducing brute luck inequality, and securing social equality. Under some circumstances, the two egalitarian criteria will align. Insofar as the benefits of an increase in the size of the mandated package come to the least-well-off and its costs are principally borne by the better off, both brute luck egalitarianism and social egalitarianism will tend to favor this expansion. We should note, however, that the two egalitarian views can also diverge. For social egalitarianism assigns particular importance to the impact of health on people's ability to perform the role of a citizen. The conditions that threaten this ability, such as disfigurement or severe

impairments in mobility, communication, or mental health, may therefore be given greater weight than their well-being impact alone would warrant (Hausman 2015, chap. 13). Social egalitarianism can therefore be expected to prompt a shift in the mandatory package of insurance towards interventions that secure capacities that are central to acting as a free and equal citizens.

Conclusion

I have put forward a revised liberal egalitarian account of mandatory health insurance. On this view, a mandate can be justified because of its role in overcoming two problems with voluntary insurance. The first is adverse selection, which occurs in a particularly problematic form when we require, on grounds of equity, that those with unfairly elevated health risks pay similar premiums to those at low risk. The second is underprovision of the public good of social equality.

I have also argued that the content of the mandatory package should not be what a representative prudent individual would purchase for themselves. Instead, it should be determined by an appeal to the following values: improving people's well-being, reducing brute luck inequality, and securing social equality. Compared to a package that a representative prudent individual would want, these values will typically prompt us to include more coverage for interventions that help the worst off, as well as more coverage for interventions that prevent the ills of domination, marginalization, and loss of self-respect.

What, then, should we say to Mark and Lucinda, sensible people of average means and moderate current health risk who are compelled to purchase a more comprehensive and therefore more costly package than they desire?

First, that in order to offer fairly-priced insurance to everyone, including people who, through bad luck alone, face higher health risks than they do now, we have to require everyone to purchase this more comprehensive package. While this will raise their current premiums, it will also lower their premiums should they eventually face high health risks. By protecting others, they are therefore also protecting themselves.

Second, we should ask them to consider that, because of its universal adoption, the package will shape the society they live in. Additional coverage for conditions that leave people especially vulnerable and isolated will result in less corrosive inequality and better prospects for people to participate in public life with dignity. In short, we should say that a mandate is necessary because fair access to health insurance and a less divided society are things that we can only produce together.

References

Arneson, Richard (1989). "Equality and Equal Opportunity for Welfare," *Philosophical Studies* 56: 77—93.

Akerlof, George (1970). "The Market for Lemons: Quality, Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics* 84: 488—500.

Anderson, Elizabeth (1999). "What Is the Point of Equality?" *Ethics* 109: 287–337.

Arrow, Kenneth (1963). "Uncertainty and the Welfare Economics of Medical Care," *The American Economic Review* 53: 941—73.

Barro, Josh (2010). "A 'Scaled-Back' Health Bill Won't Work." *Real Clear Markets Blog*. http://www.realclearmarkets.com/articles/2010/01/26/a_scaled_back_health_bill_wont_work_97609.html [Accessed January 2017.]

Bou-Habib, Paul (2006). "Compulsory Insurance without Paternalism," *Utilitas* 18: 243—63.

Braun, S. Stewart (2012). "Mandatory Health Insurance, Liberalism, and Freedom," *Public Affairs Quarterly* 26: 179—97.

Buchak, Lara (2013). *Risk and Rationality*. (Oxford: Oxford University Press.)

Buchanan, Allen (1984). "The Right to a Decent Minimum of Health Care," *Philosophy & Public Affairs* 13: 55—78.

Carrin, Guy and Chris James (2005). *Social Health Insurance: Key Factors Affecting the Transition towards Universal Coverage*. (Geneva: World Health Organization.)

Chase, Elaine and Grace Bantebya-Kyomuhendo, eds. (2015). *Poverty and Shame: Global Experiences*. (New York: Oxford University Press.)

Claassen, Rutger (2013). "Public Goods, Mutual Benefits, and Majority Rule," *Journal of Social Philosophy* 44: 270—90.

Cohen, Gerald Allen (1977). "Robert Nozick and Wilt Chamberlain: How Patterns Preserve Liberty," *Erkenntnis* 11: 5—23.

Cotlear, Daniel, Somil Nagpal, Owen Smith, Ajay Tandon, and Rafael Cortez (2015). *Going Universal: How 24 Developing Countries Are Implementing Universal Health Coverage Reforms from the Bottom Up.* (Washington, DC: The World Bank.)

Daniels, Norman (2013). "Justice and Access to Health Care," *Stanford Encyclopedia of Philosophy* http://plato.stanford.edu/entries/justice-healthcareaccess/ [accessed August 2016].

Deaton, Angus (2013). "What Does Empirical Evidence Tell Us about the Injustice of Health Inequalities?" in *Inequalities in Health*, Nir Eyal, Samia Hurst, Ole Frithjof Norheim, and Daniel Wikler (eds.). (Oxford: Oxford University Press): 263—81.

Dworkin, Ronald (1994). "Will Clinton's Plan Be Fair?" *New York Review of Books* http://www.nybooks.com/articles/1994/01/13/will-clintons-plan-be-fair/ [Accessed May 2016].

Dworkin, Ronald (2000). *Sovereign Virtue: The Theory and Practice of Equality*. (Cambridge, MA: Harvard University Press.)

Dworkin, Roland (2002). "Sovereign Virtue Revisited," *Ethics* 113: 106—43.

Dworkin, Ronald (2012a). "Why the Health Care Challenge Is Wrong," *New York Review of Books* April 2 issue http://www.nybooks.com/daily/2012/04/02/why-health-care-challenge-is-wrong/ [Accessed May 2016.]

Dworkin, Ronald (2012b). "A Bigger Victory than We Knew," *New York Review of Books*, August 16 issue http://www.nybooks.com/articles/2012/08/16/bigger-victory-we-knew/ [Accessed January 2017.]

Fourie, Carina, Fabian Schuppert, and Ivo Wallimann-Helmer (eds.) (2015). S*ocial Equality: On What It Means to Be Equals.* (New York: Oxford University Press.)

Gauthier, David (1963). *Practical Reasoning*. Oxford: Clarendon Press.

Gibbard, Allan (1984). "Health Care and the Prospective Pareto Principle," *Ethics* 94: 261—82.

Hausman, Daniel (2015). *Valuing Health: Well-Being, Freedom and Suffering.* (Oxford: Oxford University Press.)

Kurtulmus, Ahmet Faik (2012). "Dworkin's Prudent Insurance Ideal: Two Revisions," *Journal of Medical Ethics* 38: 243—46.

Lagomarsino, Gina, Alice Garabrant, Atikah Adyas, Richard Muga, and Nathaniel Otoo (2012). "Moving towards Universal Health Coverage: Health Insurance Reforms in Nine Developing Countries in Africa and Asia," *The Lancet* 380 (9845): 933—43.

Lange, Benjamin (2016). "Restricted Prioritarianism or Competing Claims?" *Utilitas*, online early, http://doi.org/10.1017/S0953820816000182

Lippert-Rasmussen, Kasper (2001). "Egalitarianism, Option Luck, and Responsibility," *Ethics* 111: 548—79.

Macleod, Colin (1998). *Liberalism, Justice and Markets: A Critique of Liberal Equality*. (Oxford: Oxford University Press.)

Mattila, Mikko, Peter Söderlund, Hanna Wass, and Lauri Rapeli (2013). "Healthy voting: The effect of self-reported health on turnout in 30 countries," *Electoral Studies* 32: 886–91.

Menzel, Paul T. (2012). "Justice and Fairness: a Critical Element in US Health System Reform." *The Journal of Law, Medicine & Ethics* 40: 582—97.

Moncrieff, Abigail (2013). "The Individual Mandate as Healthcare Regulation: What the Obama Administration Should Have said in NFIB v. Sebelius," *American Journal of Law and Medicine* 39: 539—72.

Nagel, Thomas (1970). *The Possibility of Altruism.* (Princeton: Princeton University Press.)

O'Neill, Martin (2008). "What Should Egalitarians Believe?" *Philosophy & Public Affairs* 36: 119–56.

Olsaretti, Serena (2009). "Responsibility and the Consequences of Choice," *Proceedings of the Aristotelian Society* CIX, Part 2: 165—88.

Otsuka, Michael (2012). "Prioritarianism and the Separateness of Persons," *Utilitas* 24: 365–80.

Otsuka, Michael and Alex Voorhoeve (2009). "Why It Matters that Some Are Worse Off than Others," *Philosophy and Public Affairs* 37: 171—99.

Otsuka, Michael and Alex Voorhoeve (forthcoming). "Equality versus Priority," in S. Olsaretti (ed.) *The Oxford Handbook of Distributive Justice*. (Oxford: Oxford University Press.)

Ottersen, Trygve, Reidun Førde, Meetali Kakad, Alice Kjellevold, Hans Olav Melberg, Atle Moen, Ånen Ringard, and Ole Frithjof Norheim (2016). "A New Proposal for Priority Setting in Norway: Open and Fair," *Health Policy* 120: 246—51.

Quiggin, John (1993). *Generalized Expected Utility Theory. The Rank-Dependent Model.* (Boston: Kluwer Academic Publishers.)

Rawls, John (1999). *A Theory of Justice, revised 2nd edition*. (Oxford: Oxford University Press.)

Rawls, John (2001). *Justice as Fairness: A Restatement.* E. Kelly (ed.). (Cambridge, MA: Harvard University Press.)

Rulli, Tina and David Wendler (2016). "The Duty to Take Rescue Precautions," *Journal of Applied Philosophy* 33: 240—58.

Satz, Debra (2010). *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. (Oxford: Oxford University Press.)

Segall, Shlomi (2010). *Health, Luck and Justice*. (Princeton: Princeton University Press.)

Smets, Kaat and Carolien van Ham (2013). "The Embarrassment of Riches? A Meta-Analysis of Individual-Level Research on Voter Turnout," *Electoral Studies* 32 (2): 344–59.

Smith, Adam (1982 [1790]). *The Theory of Moral Sentiments*, D. D. Raphael and A. L. Macfie (eds.). (Indianapolis: Liberty Fund.)

Sreenivasan, Gopal (2014). "Justice, Inequality and Health," *Stanford Encyclopedia of Philosophy* https://plato.stanford.edu/entries/justice-inequality-health/ [Accessed January 2017].

Stemplowska, Zofia (2009). "Making Justice Sensitive to Responsibility," *Political Studies* 57: 237—59.

Tawney, Richard (1964). *Equality* (5th edition). (London: Unwin.)

Temkin, Larry (2001). "Inequality: A Complex, Individualistic, and Comparative Notion," *Philosophical Issues* 11: 327—53.

Vallentyne, Peter (2002). "Brute Luck, Option Luck, and Equality of Initial Opportunities," *Ethics* 112: 529–57.

Voorhoeve, Alex and Marc Fleurbaey (2016). "Priority or Equality for Possible People?" *Ethics* 126 (2016): 929—54.

WHO (2014). *Making Fair Choices on the Path to Universal Health Coverage.* (Geneva: World

Health Organization.)

Appendix: Why the RPIT fails even when paired with an unorthodox decision theory

The argument in Section 2 assumes a measure of well-being that relies on orthodox decision theory. A familiar objection to this theory (and my related assumption that prudence requires maximizing expected well-being) is that when one is evaluating a prospect from a prudential point of view, one should be permitted to be especially averse to worse outcomes. A well-known alternative to the orthodoxy, known as "rank-dependent" decision theory, indeed allows one to assign a lower-ranked outcome more weight than its probability alone would warrant (Quiggin 1993; Buchak 2013).

It is noteworthy that the fundamental objections to the RPIT would remain even if one were to employ this rank-dependent decision theory. To see why, I shall consider a pair of contrasting cases. Let us start with what I will call the Equal Health Case. In this scenario, society consists entirely of people who are all "in the same boat," in the sense that they will either all end up Ill or all end up Well and are equally likely to end up in either situation. The government can choose between policy A and B. Policy A does relatively little for those who are sick; it therefore leaves Ill much worse off than Well. Policy B, by contrast, ensures complete equality between Ill and Well. However, it comes at a cost c in expected well-being. Table A1 describes this case. Note that the columns describe states of the world. Assume that $0 < c < 20$, so that, compared to A, B always improves people's situation if they fall ill.

Table A1. Well-being outcomes in the Equal Health Case

State of the world Policy	Everyone Ill (probability = 0.5)	Everyone Well (probability = 0.5)
A	40	80
B	60 - c	60 - c

Now, on standard decision theory, A is prudentially best for each person, because it maximizes each person's expected well-being. However, on rank-dependent decision theory, it is permissible to give some extra weight to the worst possible situation, so that for some c, B is prudentially better for each person. Let c^{Equal} be the largest cost for which one could, on prudential grounds alone, choose B in the Equal Health Case. The rank-dependent RPIT will therefore regard A as permissible for all cases in which B is at least as costly as this, that is, for all $c \geq c^{Equal}$. Indeed, it will require A for all cases in which B is more costly than this, that is, for all $c > c^{Equal}$.

Next, consider the Unequal Health Case, in which there is no risk. Half the population is Ill, the other half is Well. The government must decide between policy A, which leaves the Ill much less well off than the Well, and B, which generates complete equality, again at cost c, with $0 < c < 20$. This case is represented in Table A2. Note that the columns are now segments of the population.

Table A2. Well-being outcomes in the Unequal Health Case

People Policy	Ill (half of the population)	Well (half of the population)
A	40	80
B	60 - c	60 - c

It should be obvious that there are important differences between the Equal and Unequal Health Cases. First, if one chooses A in the Equal Health Case, there is no inequality. By contrast, if one chooses A in the Unequal Health Case, there will be substantial unfair inequality in both prospects and outcomes. Second, the Equal Health Case involves only a trade-off between two different possible futures of each individual. By contrast, the Unequal Health Case involves a trade-off between the interests of separate people. I submit that these differences make it more important to improve the worst possible situation in the Unequal Health Case than in the Equal Health Case. It follows that there is some c for which it is permissible to choose A in the Equal Health Case but impermissible to choose A in the Unequal Health Case, because one ought to do more to improve the situation of the ill in the latter.

However, the rank-dependent RPIT cannot recognize this. For in the Unequal Health Case, the RPIT will tell the government to choose the package that it would be prudent to purchase on one's own behalf if one had average risks, that is, were equally likely to end up Ill or Well. It follows that the RPIT would draw the same conclusion in the Unequal Health Case as in the Equal Health Case: regard A as permissible for all $c \geq c^{Equal}$ and as required for all $c > c^{Equal}$. The RPIT is, in this sense, wholly insensitive to the presence or absence of unfair

inequality and to whether the trade-offs in question are between two potential futures of the same individual or instead between the interests of separate individuals.

Of course, neither of these cases represents a realistic scenario. But each of these cases represents, in an isolated form, one aspect that is present in real-world cases of universal, mandatory purchase of an insurance package. The RPIT's failure to respond appropriately to the differences between these cases therefore generates a problem for this principle in real-world scenarios that consist of a mixture of the elements that each case models in isolation. The RPIT, whether paired with orthodox or rank-dependent decision theory, will recommend exactly the same policy in all scenarios, no matter whether they are closer to the ideal type represented by the Equal Health Case or closer to the ideal type represented by the Unequal Health Case. It is therefore also unresponsive to inequality and to people's competing claims under realistic circumstances.

---

[2] Gibbard (1984) makes a similar argument on utilitarian grounds.

[3] While it differs from their individual proposals in some details, this formulation of and motivation for the RPIT picks out the relevant common elements in proposals by Gibbard, Dworkin, Menzel, and Kurtulmus. The

formulation employed avoids some of the problems with Dworkin's (1994, 2000) version of the RPIT highlighted in Macleod (1998, pp. 92-6) and Kurtulmus (2012).

[4] Of course, in reality, no adult enjoys such certainty, but we grant Healthy knowledge of her unusual fortune for the sake of a simplifying idealization of the situation of individuals who are known to have better than average health prospects.

[5] In the Appendix, I argue that one ought to reject RPIT even if one pairs it with a well-known rival decision theory.

[6] I follow Gibbard (1984) in assuming this measure, which draws on idealized preferences which respect the von Neumann-Morgenstern axioms. This measure does not presuppose any particular view on what well-being consists in. One might maintain both that well-being consists in something other than preference satisfaction and that the specified idealized preferences fully track the magnitude of this other thing (Otsuka and Voorhoeve 2009, pp. 172-3n3).

[7] While this second objection is related to the first, it is, as Otsuka (2012) shows, logically distinct. See Lange (2016) for further discussion of the way in which being worse off than others affects the strength of a person's claim.

[8] See Voorhoeve and Fleurbaey (2016) and Otsuka and Voorhoeve (forthcoming) for defenses of a pluralist egalitarian view and Ottersen et al. (2016) for a proposal to use such a view to determine health services provision in Norway.

[9] There are well-known difficulties in adequately spelling out the distinction between brute luck and option luck (Vallentyne 2002, pp. 531-4). There are also powerful arguments that even inequalities due to differential option luck are objectionable because people cannot reasonably be held substantively responsible for them (Lippert-Rasmussen 2001). I shall not engage with these challenges, since I argue in Section 4 that even an egalitarian who is generally willing to accept inequalities due to differential option luck should object to such inequalities in this case.

[10] My argument that a focus only on the individual impacts of a transaction omits consideration of its distributional and relational impacts has parallels to Cohen's (1977) critique of Nozick's (1974) "Wilt Chamberlain argument."

[11] Since the view focuses on citizens, a question arises what it requires with respect to the treatment of resident non-citizens. For simplicity, I set aside this issue.

[12] It may be useful to compare my proposal with the view that mandatory basic health insurance contributes to the public good of living in a society that meets its humanitarian obligations towards its members (Buchanan 1984; Segall 2010, pp. 78-80; Menzel 2012.) It is, plausibly, a precondition for a society of equals that, out of appreciation for the value of each person's life, society is committed to meeting citizens' basic needs when this can be done at reasonable cost (cf. Rawls 1999, p. 298). However, such a commitment is not sufficient for social equality. The social egalitarian public goods that I discuss here are therefore more encompassing than this humanitarian public good.

[13] Depending on citizens' attitudes, this collective action problem may therefore take one of two forms. One is a multi-person, one-shot Prisoner's Dilemma. This occurs if, in the situation outlined in Table 3, each citizen rationally chooses Medium no matter what others do. But citizens may also be conditional cooperators. In Table 3, such a citizen will prefer Medium if he believes others will choose Medium, and Large if he believes others will choose Large. If all citizens are conditional cooperators, then their interaction has the structure of an Assurance Game. The purpose of a mandate is different in each case. If the problem is a Prisoner's Dilemma, then it is to force cooperation on a social optimum; if it is an Assurance Game, then it is to generate coordination on a social optimum by providing assurance that others will contribute. I am grateful to Joshua Cohen for pressing me to clarify this point.

[14] Indeed, Dworkin (2000, pp. 155-8) endorses a "principle of correction" that requires that the state prevent such market failures. It follows that, if Dworkin had accepted my argument that the social egalitarian goods produced by health insurance are public goods, then he would have had to withdraw his claim (quoted at the start of Section 3) that a society of Avys who were all left free to choose Medium "could not be criticized on grounds of justice." I should note that Dworkin (2002, p. 114) appeals to his principle of correction as a justification for mandatory "personal accident" insurance on the grounds that an individual's decision not to purchase such insurance imposes costs on "the rest of the community, including employers and dependents." Dworkin does not clarify the costs he has in mind, but perhaps he was thinking of the emergency medical care that hospitals are, out of humanity, obligated to provide in many countries (e.g. under the Emergency Medical Treatment and Labor Act in the U.S.), as well as public welfare support for the person's dependents. If so, then this argument of Dworkin's for mandatory personal accident insurance would rely on an appeal to the fair sharing of the costs of meeting humanitarian obligations. As mentioned in n. 12, meeting these obligations is

but one part of establishing social equality. Dworkin therefore seems to overlook the contribution of health

insurance to the social egalitarian public goods that I discuss here.