Public value, maximisation and health policy: an examination of Hausman's Restricted Consequentialism

Article type: Original article

Manuscript title: Public value, maximisation and health policy: an examination of Hausman's Restricted Consequentialism

Running title: Public value, maximisation and health policy

Author: James Wilson, Department of Philosophy, University College London

Address for correspondence: Dr James Wilson, Department of Philosophy, University College London, Gower Street, London WC1E 6BT. Tel: 0207 679 0213. Email: james.wilson@ucl.ac.uk

Summary: In the book Valuing Health, Daniel Hausman sets out a normative framework for assessing social policy, which he calls restricted consequentialism. For the restricted consequentialist, government policy-making not only is, but ought to be, largely siloed in individual government departments. Each department has its own goal linked to a fundamental public value, which it should pursue in a maximising way (subject to constraint by other non-consequentialist values). I argue that Hausman's argument appears to be internally inconsistent: his case for thinking that health policy should default to a form of maximisation is plausible only if a much narrower vision of the goals of policy is adopted than Hausman thinks appropriate in the case of education. Second, it turns out that none of Hausman's analysis helps us with the crucial question of how maximisation should be constrained by other values -- a question that even on Hausman's account looks to be crucial, and that will be even more important if adopt a broader perspective on the purposes of health policy than Hausman allows.

Key words: Accountability for reasonableness, Muddling through, Restricted Consequentialism, Social Policy.

Acknowledgements: Thanks to the audience at the Valuing Health conference at UCL where this paper was initially present.

Conflicts of Interest: I have no conflict of interest.

Public value, maximisation and health policy: an examination of Hausman's Restricted Consequentialism

Introduction

Daniel Hausman's principal aim in *Valuing Health* (2015) is to determine how health should be valued for the purposes of public policy.[1] Hausman also has a more ambitious agenda, which emerges especially in the final two chapters of the book, namely to advance an approach to setting priorities in public policy more generally. This is *restricted consequentialism,* which Hausman thinks is exemplified by his approach to health policy, but is also relevant to many other areas of public policy.

Hausman's argument for restricted consequentialism makes four main moves, each of which is analysed in more detail below.

1. An argument that liberal governments should aim to promote public rather than private values.

2. An argument about the need for a core of maximising consequentialism in social policy.

3. An argument about causal complexity and the availability of information, which entails that governments should (as they in fact do) pursue a number of different sub-goals such as health or security without aiming to focus too much on an overall goal such as well-being, even if success in promotion of well-being should be the goal by which governments are judged.

4. An argument about the kinds of goods (such as solidarity, compassion and fairness) that should constrain maximisation.

Restricted consequentialism holds that "a policy that satisfies relevant normative constraints is acceptable if and only if no alternative results in better consequences". (p. 226) For the restricted consequentialist, government policy-making not only is, but ought to be, largely siloed in individual government departments, each of which has its own goal linked to a fundamental public value. Each department should pursue its respective goals in a maximising way (subject to constraint by other non-consequentialist values).

1 All page numbers are to this work unless otherwise stated.

My aim here is to examine the argumentative steps by which Hausman reaches restricted consequentialism. As we shall see, close analysis shows a number of problems for restricted consequentialism—problems that also threaten Hausman's approach to health policy. Two challenges stand out. First, Hausman's argument appears to be internally inconsistent: his case for thinking that health policy should default to a form of maximisation is plausible only if a much narrower vision of the goals of policy is adopted than Hausman thinks appropriate in the case of education. Second, it turns out that none of Hausman's analysis helps us with the crucial question of *how* maximisation should be constrained by other values – a question that even on Hausman's account looks to be crucial, and that will be even more important if (as I shall argue) we need to adopt a broader perspective on the purposes of health policy than Hausman allows.

Understanding restricted consequentialism

Hausman distinguishes between the private and the public value of health. The private value of a health state is its "contribution to whatever the individual cares about or should care about" (p. 158). Hausman does not seem to give an explicit definition of public value, rather saying that "determining the public value of a health state is a task for political philosophy whose outcomes depend upon what the aims of public policy ought to be" (p. 154).

Hausman grounds his account of public value in a specific understanding of the role of the liberal state, recognising that he does so more by way of stipulation than by argument.[2] On Hausman's liberal view, the state's goals should be to "create an environment that secures the basic prerequisites for common activities and competencies, including especially, the competencies required by citizenship and the ability to formulate and pursue personal interests within the limits set by the rights of others." (p. 159) This entails a clear distinction between private and public values: "public values are not an aggregation of personal values, because the concerns of a liberal state are distinct from the objectives of its citizens, among which it should be largely neutral." (p. 159)

2 At the end of chapter 12 Hausman acknowledges that his public value account of health will be useful only for liberal states. He does not address the question of whether a different account of public value could be devised for some non-liberal states. We could distinguish two different possibilities. On the first interpretation, public value is not specifically liberal: the public value of a health state (or educational state) is its contribution to whatever the state cares about or should care about. On this conception of public value, activity that ensured religious observance might have public value in a religious state. On the second reading of public value, public value is a specifically liberal concept, which is connected up to the idea of *public justification*: public values are those that are justifiable on the basis of a certain kind of reasoning. (Vallier and D'Agostino, 2014). Given that Hausman explicitly restricts himself to liberal societies, his account of public value does not require him to choose between these options.

Hausman's reason for thinking that approaches to public policy should have a core of maximisation cannot be that consequentialism just is the correct moral theory, given that his aim is to defend an approach in which maximising approaches to public value are constrained by a range of nonconsequentialist values. Indeed, it seems that Hausman thinks that restricted consequentialism should be favoured as an approach to public policy even by those who do not adopt consequentialism at the level of fundamental moral theory.

So why assume that social policy requires a base of maximising consequentialism? At times, Hausman seems to suggest that the centrality of maximising measured benefits in social policy is simply a given from which we should start: "[i]n our era, policy analysts favor maximizing quantitatively measured benefits whenever it is feasible and ethically permissible to do so." (p. 224) But even if true, the claim would accurately relate a sociological fact about policy makers, rather than a normative or practical requirement on policy making. Why couldn't there be an approach to social policy that did not prioritise maximising quantitatively measured benefits?

Here, I venture a potential reply on Hausman's behalf, namely that government accountability to citizens *requires* quantifying the benefits and costs of policies. This answer is implicit in Hausman's text and in his reference (p. 218) to the work of Norman Daniels (2007) on accountability for reasonableness, but it does not come through as clearly as it might. Here is a reconstruction of such an argument from accountability.

The argument starts from the duty on governments to be accountable to their citizens for the way that power is exercised. Given that there are competing priorities for government attention, and that governments cannot pursue all the policies that might benefit at least some individuals, funding any particular policy will have the implication that other policies are not pursued. In other words, governments need to take into account opportunity costs.

Citizens have a general interest in tax money being well spent; but they also have a more specific interest where they stand to lose as individuals from a decision such as to close a public library, or where the decision will subject them to coercion if they fail to comply (for example, increasing taxes). These interests of citizens qua citizens create a duty of accountability.

On this view, if a government cannot give a plausible account of why one policy is being pursued rather than another, then it is failing in its duty of accountability. In the case of healthcare rationing, Daniels argues that there is a duty to show that the budget has been used in a way that is a reasonable use of resources, given the relevant constraints.[3] We can think of governments as under similar duties for other departments beyond health too.

3 "*Relevance Condition*: The rationales for limit-setting decisions should aim to provide a reasonable explanation of how the organization seeks to provide value for money in meeting the varied health needs of a defined population under reasonable resource constraints." (Daniels, 2007, p. 118)

Setting a goal for the government department that can be measured with a single cardinal number, and which the government department can attempt to maximise, provides one way of discharging the generalised requirement for accountability. If the government does so, it will be able to provide *an answer* to the question of why it has allocated resources in one way rather than another. If a government can show that no alternative policy results in better consequences in terms of the agreed goal, then this would be "accepted as relevant by (fair minded) people who are disposed to finding mutually justifiable terms of cooperation". (Daniels, 2007, p. 118)

However, even such an accountability based justification for maximising public value faces severe challenges. As we shall see below, Hausman himself argues that in certain areas of government policy, maximisation of public value is not possible. Where maximisation of public value is not possible, it cannot be required by duties governments bear to their citizens.

Are there scalar measures of public value?

Hausman is commendably honest in revealing that—at least in some policy areas—it seems implausible to think that maximising public value would be a sensible policy goal. In particular, he suggests that the values involved in an education system from a public perspective "may be too diverse" to be summed up in a single number:

> What is wanted of an educational system from a public perspective includes cultivations of the knowledge, skills, and abilities needed to thrive as a family member, friend, worker, and citizen. Refining and weighing these is problematic. I do not know if it is possible to define an informative scalar measure of the public value of education with which to summarize how well an educational system is functioning, and I'm skeptical about whether it is possible to acquire useful information concerning the cost effectiveness of educational policies. (p. 228)

It would seem that on Hausman's view there may be some things that can be maximised within education, such as the percentage of children who get secondary education credits at grades A-C (as English schools attempt to do), but we should not confuse achieving such goals with maximising the public value of education.[4] Either we can have a single number measure of educational value that does not adequately track public value, or we can have a measure of public value that cannot be summed up in a single value or is not informative informative for public

4 This is not to suggest that those who do promote such educational policies think that exam results are an adequate measure of the public value of education. It seems more likely that they think exam results are proxies for something more important. Thanks to an anonymous referee for prompting me to be clearer here.

policy purposes. In any case, maximisation of the public value of education seems to be off the table.

Hausman's account of the diversity of values involved in education seems plausible. But as soon as it is acknowledged that the values involved in *education* are too diverse to admit of a scalar measure of the public value of education, then his account faces an obvious question: why aren't the values involved in other policy areas such as *health* also too diverse for there to be a useful single number measure of *their* public value?

In chapter 13, Hausman gives us a measure of the public value of health that takes into account only two dimensions: distress and activity limitation. But this account is very narrow when compared to the yardstick he used for the public value for education, namely "how well an educational system is functioning". Parity would demand that the relevant yardstick for the public value of health should be how well a health care system is functioning, rather than how much certain health states are improved in respect of their public value.

This is an important point, as the values that should animate a health care system are commonly thought to go well beyond maximising health states. If we ask what is relevant to how well a health care system is performing, then it seems that there are a number of relevant goals or values that are not captured by Hausman's limitation/distress model. Here are a few:

- Providing a comprehensive service free at the point of use.
- Candour in admitting mistakes. (Francis 2013)
- Avoiding catastrophic out-of-pocket health expenditures.
- Respect for professional standards.[5]

It may be that Hausman would wish to argue that each of these cases is best thought of *not* as a goal or value to be promoted, but rather as a constraint on the maximisation of the public value of health as he conceives it. Certainly some values will be plausibly construed in this way: candour for example. The hope is that the health care system will not adjust its candour about its mistakes to whatever will maximise the amount of health, but rather treat "ensuring that patients harmed by a healthcare service are informed of the fact and that an appropriate remedy is offered, whether or not a complaint has been made or a question asked about it" (Francis 2013) as a constraint on maximisation.

5 This was a factor in criticism of NICE's initial decision to allow age related macular degeneration patients to go blind in their weaker eye before the drug lucentis could be administered. The implications of this case are discussed at greater length in Weale et al. (2014)

Other cases are less easily assimilated. For example, providing a comprehensive health service that is free at the point of use is taken by many in the UK *to be* what good health care system should aim at. From this perspective, providing a comprehensive service that is free at the point of use is not a *constraint* on the maximisation of the public value of health, but rather the goal of the health care system as a whole. On such a view, a cost effectiveness threshold such as the one the National Institute for Health and Care Excellence (NICE) operates is a side effect of trying to maintain a comprehensive service on the basis of a fixed and limited budget, rather than the result of an attempt to maximise health outcomes.[6]

Hausman does give an extended role for public values such as opportunity, compassion, fairness, solidarity and equal respect (p. 215) as constraints on maximisation. So he might think that these values could provide a defence against the charge that his account of the public value of health is problematically narrow. But appeal to these values as constraints on maximisation does nothing to rebut the claim that it is reasonable to think that e.g. reducing out-of-pocket payments is an important and publicly justifiable goal of a healthcare system that is separate from and irreducible to maximising health gain.[7]

The problem, in brief is this: Hausman argues that we should shift from measuring the private to the public value of health, because the private value that health states have for individuals is not relevant from the perspective of the liberal state. However he assumes that in the shift from private to public value the goal of health policy will not change fundamentally: the shift will be from constrained maximisation of the private value of health states to constrained maximisation of the public value of health states. But as is apparent from Hausman's own reflections on education, there is no reason to conceive of public value so narrowly. It makes at least as much sense to think about public value in terms of the features of the systems by which governments seek to promote and protect goods such as health or education. If we adopt this wider perspective,

6 Rumbold et al. (2016) argue that it is more accurate to view the cost effectiveness procedures NICE actually adopts as a form of constrained satisficing, rather than one of constrained maximisation. This is because no attempt is made within the system to maximise the amount of QALYs produced: rather the system operates in such a way that interventions that cost more than £20,000–£30,000 per QALY will not usually be funded unless there are certain features—such as severity or life extension at the end of life—that warrant going beyond the usual cost effectiveness thresholds.

7 For such an approach, see for example World Health Organisation (2014). Moreover, it is unclear how Hausman thinks these constraining values are justified: are they supposed to be derived from the nature of liberalism, or are they are supposed to be values that are compelling regardless of whether one is a liberal or not? Either way, questions arise for Hausman: if all these values *are* related to the core of liberalism, then it would seem that these are over-arching values that should play a structural role in policy everywhere. Perhaps policy should then not be siloed in the way that Hausman thinks it should be; perhaps we will need to ask different kinds of question: for example, what does the kind of solidarity we owe one another as human beings look like in different domains such as health and education? But if it is legitimate to appeal to non-liberal values in constraining the pursuit of liberal public value, then why place so much emphasis on defining public value via the idea of liberalism in the first place?

it is by no means clear that there could be a single scalar public value for a particular area of government activity.

Overall, it seems relatively uncontroversial that government accountability requires some kind of way of transparently prioritising different policies; but it is not at all clear that duties of accountability create any kind of requirement for there to be maximisation in order to present reasons that fair minded people will accept. While some form of constrained maximisation is a very common approach in public policy, it is not the only game in town.

Public value, complexity and information availability

While seeming to favour constrained maximising approaches to public value at a theoretical level, Hausman argues that in practice the ability to gather the information required to adopt a maximising policy is limited. As the massive literature on the social determinants of health (among other literatures) have shown, health (and health policies) are deeply causally entangled with policies in other areas: "Health policies affect educational outcomes, and educational policies, like environmental policies, transportation policies, and the criminal law affect population health" (p.224). In fact, things are much worse than this: it's not just that health and education are intertwined in a top level way; that educational outcomes affect health outcomes and vice versa; this entanglement goes all the way down, and is key to why even small scale policies succeed or fail in practice.[8]

In the face of these kinds of challenges, Hausman points out that contemporary governments attempt to simplify the problem space: rather than defining their goal as "the promotion of well-being by manipulating the particular causal factors within their purview, contemporary governments assign different goals to different sectors. They do this because there is no feasible alternative, even if the ultimate concerns are well-being and fairness." (p. 197) On Hausman's view, policy makers cannot reasonably expect health economists to be able to estimate the effect of health policies on well-being, but they can reasonably expect health economists to be able to estimate the effect of health policies on health. (p. 197) Therefore policy makers should adopt a siloed approach, which focuses mainly on the distribution of goods such as health that are narrower in scope than well-being.[9]

8 See for example Cartwright and Hardie (2012). I examined these problems in Wilson (2009).

9 Moreover, on the basis of these information problems, Hausman argues for incrementalism: "When ambitious policies are proposed, it is tough to anticipate the benefits and drawbacks that may be obvious once policies are implemented. Although tinkering with existing policies may lock policy makers into institutions that are far from optimal, it may also be the only feasible way to avoid disastrous mistakes." (p. 226) It is not clear to me what role the idea of maximisation is still playing in an explicitly incrementalist approach. Indeed, in introducing such a "muddling through" approach Lindblom (1959) took himself to be explicitly opposing maximising approaches.

Hausman admits that "this division of labour and separation of objectives has some unfortunate consequences", as "different agencies may institute policies that duplicate or undercut one another" (p. 197). But despite these risks, he argues that there is no feasible alternative, even for those who think that promotion of well-being should be the overall goal of governments.

This seems plausible on the surface, but rather more suspect on closer analysis. First, decisions about budgets for different departments are obviously impossible to make on the basis of the separate goals for each department (unless we have a way of trading off health and education against one another, how can we decide what the relative size of the health and education budgets should be?) So it seems that we need to be able to make judgements about the well-being effects of policies – even if at a fairly rough and ready level, in order to be able to guide budget allocation decisions between government departments.

Second, the fact that one is restricting one's gaze to health makes the problem of measurement somewhat easier than focusing more broadly on well-being, but it would be wrong to think that narrowly focused policy questions are always or nearly always tractable, whilst policy questions focused on well-being are usually intractable. The health effects of large scale decisions within one department (such as whether to merge the health and social care budgets, or de-institutionalisation through "care in the community", or split a previously monolithic health care system into purchasers and providers) seem to be much more difficult to estimate than the cross sector effects of creating a new bus route.

If Hausman's argument for departmental silos is epistemic, then it should be responsive to the different relative states of knowledge: it will be easier to estimate the effect of some policies on overall well-being than it is to estimate the effects of others on health alone. Moreover, the availability of better data should make it more tractable to think about the cross sector effects of policies: so over time, and with better cross sector working we may come to be able to feasibly pursue well-being as a policy goal.

However, it may be that normative claims about the unfeasibility of a single scalar measure of the public value of well-being are also playing a part in the background. Over the course of the book, Hausman raises a number of difficulties for measuring the *private* value of well-being (in particular the apparent heterogeneity of well-being, and some difficulties with cost-benefit analysis); and while he does not explicitly defend the claim that it is not possible to define an informative scalar measure of the public value of well-being, it seems reasonable to suppose that this is what he thinks (given that well-being as a good is more complex than education). If so, it may be that Hausman thinks that, even given perfect information, it would not be feasible to have

such a measure of well-being, and because of this, it would be a mistake ever to take maximisation of the public value of well-being as a policy goal.

How do non-consequentialist commitments constrain maximisation?

Hausman's vision of restricted consequentialism for public policy is that policies are first "appraised in terms of their promotion of one among an array of the social goods, rather than in terms of an overall objective" (p.226), and second that maximisation of the relevant social good is constrained by non-consequentialist commitments. Restricted consequentialism gives us a structure for policy evaluation, but this structure needs to be specified in order to be determinate enough to guide policy. As Hausman points out: "one needs to define the normative constraints and to specify what it is for consequences to be 'better'" (p.226)

Hausman says little about how the non-maximising elements (values like solidarity, and compassion) should bear on the maximising elements. It is one thing to say that the non-maximising elements "constrain" maximisation, but what we really need to know is how this occurs. It's clear that on Hausman's account these non-maximising values do not operate as anything like absolute or near absolute side-constraints, but little positive is said about how they do operate.

This seems to be a crucial omission. The vast majority of *Valuing Health* is taken up with arguing about how to value health. By chapter 13, we get to the position that we should be focusing on the public value of health rather than its private value; but when we get to the discussion of the ethical challenges of healthcare resource allocation in chapter 16, Hausman points out that the most discussed ethical problems in healthcare resource allocation ("fair chances", "discrimination", "severity" and "non-aggregation") are not problems with the *measure* of the effectiveness of health policies, but rather with the use of cost-effectiveness itself. Therefore, "improving that measure does not by itself answer them." (p. 205)

I wasn't sure if Hausman thinks that shifting from private to public value will make it easier in *any* way to solve ethical problems about the use of cost-effectiveness considerations within social policy. One reason for scepticism here is that someone could think that values such as compassion or solidarity should constrain attempts to maximise health outcomes, even if they thought of the value of health in terms only of its private value.[10]

10 This is arguably the way that NICE's current social values operate.

So, overall the worry is that even if shifting to valuing health (and other goods in social policy) in terms of public rather than private value is useful for clarifying how governments should value these goods; and even if there are useful scalar measures of public value to be had, this will leave untouched the most significant ethical problems in policy making, which concern clashes between values.

Conclusion

Hausman's restricted consequentialism is an important attempt to say something fresh about an extremely difficult question: how can normative thinking best guide social policy? However, the account as it stands requires significant clarification and enhancement to be usable as the basis for public policy decisions.

There are three main moves in Hausman's account. First, that public policy should concern itself with promoting public rather than private value. Second, that the appropriate policy to adopt towards public value in public policy is maximisation. Third, that maximisation should be appropriately constrained by a range of other values.

Each of these moves has revealed significant problems. First, while public value on Hausman's account gives us a less problematic measure of the cost effectiveness of public policy than private value does, Hausman admits that by itself it does nothing to help solve what are usually thought to be the more fundamental questions of the ethics of cost effectiveness. Second, restricted consequentialism will be appropriate as an approach to policy only where there are scalar values that can usefully be maximised. However, Hausman himself doubts whether a useful scalar measure of public value could be designed in areas of government policy such as education, and we have seen that parity of reasoning would also lead us to be dubious about the possibility of a useful scalar measure of the public value of health. So, even those sympathetic to Hausman's conception of public value will judge that a commitment to public values does not entail that scalar measures of public values are possible: maximisation of public value might be the exception rather than the rule. Finally, Hausman offers only a fairly sketchy account of the values that should constrain maximisation, and the way in which they should do so.

References

Cartwright, N., & Hardie, J. (2012). *Evidence-based policy: a practical guide to doing it better*. Oxford University Press, USA.

Daniels, N. (2007). *Just Health: Meeting Health Needs Fairly*. Cambridge: Cambridge University Press.

Francis R. (2013). *Report of the Mid Staffordshire NHS Foundation Trust Public Inquiry*. London: The Stationery Office.

Hausman, D. M. (2015). *Valuing health: Well-being, freedom, and suffering*. Oxford: Oxford University Press.

Lindblom, C. E. (1959). The science of muddling throughâ€ . *Public administration review* 19(2): 79--88.

Rumbold B., Weale A., Rid A., Wilson J., and Littlejohns P. (2016). "Public Reasoning and Health Care Priority Setting: The Case of NICE". *Kennedy Institute of Ethics Journal* (in press).

Vallier, K. and D'Agostino, F. (2014). "Public Justification", *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), available at http://plato.stanford.edu/archives/spr2014/entries/justification-public/.

Weale A., Rumbold B., Rid, A., Wilson J. and Littlejohns P. (2014). "Cost-effectiveness and Social Values Recommendations For Revising Nice's Social Value Judgements: an Open Letter to the National Institute for Health and Care Excellence.", available at: www.ucl.ac.uk/socialvalues/downloads/files/Open_Letter_to_NICE

World Health Organisation (2014). *Making fair choices on the path to universal health coverage: Final report of the WHO Consultative Group on Equity and Universal Health Coverage*. Geneva: World Health Organisation.

Wilson, J. (2009). Towards a Normative Framework for Public Health Ethics and Policy. *Public Health Ethics*, 2(2):184--94.