

Welfare Economics

Antony Eagle

Choices, Models and Morals » Lecture 9

Contents

- Welfarist Policy Analysis
- Cost-Benefit Analysis

Welfarist Policy Analysis

Where we are at with welfare

- › At the end of **lecture 8** we ran into difficulties whichever way we turned in the analysis of wellbeing.
- › We wanted a theory, unlike the pure ‘objective list’ theory, that would give some connection between preference and wellbeing – a theory that allowed wellbeing-enhancement to be **intrinsically motivating** – but also allowed that people might come to have **pathological** preferences: preferences for their own oppression, preferences that don’t track what is genuinely beneficial to them.
- › Substantive theories have some prospect of making well-being something intrinsically **preferable**. But the main candidate analyses face two problems:
 1. How to **calibrate** across people? How can we measure across people the amount of pleasure or happiness a certain outcome yields?
 2. Even if we can calibrate and measure it, should we? The **hedonic monster** is someone who gets vastly more pleasure from a given outcome than anyone else, measured on a common scale. But, intuitively, that person’s pleasures shouldn’t count more than other people’s when it comes to policy in a democratic state (even if, from the God’s eye view, it would be morally better to allocate more goods to those who would get more benefit from them).

Preference and Evidence for Wellbeing

- › This concern about fairness seems to require a ‘common currency’ that cannot be distorted by hedonic monsters or the like: preferences are a natural candidate.
 - › If A and B have their preferences satisfied, that is a fact that is public and comparable between them. But since there is no scale, only a rank ordering, there is no sense in which one could get more wellbeing out of having their preference satisfied than the other – even if one is a hedonic monster and the other is not.
- › To address the problems with preferentialism as an account of what wellbeing **is**, we can opt for **no theory** of well-being – i.e., treat it as a primitive concept – but to take preference satisfaction as **defeasible evidence** for well-being (Reiss 2013: 225–27).
 - › And to take actual choice as defeasible evidence for preference (Reiss 2013: 32–35)!
- › On this view, preferences indicate, other things being equal, a person’s wellbeing; the satisfaction of those preferences is the public dimension of whether some policy conduces to their wellbeing.

A Qualitative Approach

- › There might also be a private dimension to wellbeing: e.g., what might determine your preference ranking might be something like how much pleasure you get from various options.
- › The problems of calibration and the hedonic monster militate against extending this private ground to interpersonal comparisons.
- › But if we have **intrapersonal** comparisons of wellbeing, e.g., a private ground for wellbeing as above, the following principle (making use only of individual preference) will generate some useful comparative claims about policies that apply **interpersonally**:

Principle of Personal Good (PPG) ‘If X is better for someone than Y , and Y is not better for anyone than X , then X is morally better than Y .’ (Hausman, McPherson, and Satz 2017: 151)

- › If a possible action would make no one worse off and someone better off, PPG will say that action is an improvement over the *status quo*.
 - › Sometimes this is **hedged** as holding only *ceteris paribus*: ‘an increase in one individual’s well-being is, others’ being equal, a good thing (Samuelson 1950: 11). This is then the principle of **minimal benevolence**; it says only that we have **some reason to perform** actions that improve someone’s wellbeing and decrease no-one’s.

The Pareto Principle

- › The PPG is **neutral** on the correct conception of wellbeing.
- › If we equate wellbeing with preference satisfaction – admittedly, something we objected to last time – the PPG is equivalent to:
Pareto Principle ‘If someone prefers X to Y , and no one prefers Y to X , then X is morally better than Y .’ (Hausman, McPherson, and Satz 2017: 151)
- › If we only make use of preference as a defeasible **indicator** of wellbeing, then the Pareto Principle holds only *ceteris paribus* – ‘other things being equal’.
- › This is a **sufficient condition** for one outcome being better than another. It is not a **definition** - there may be additional ways for X to be better than Y .

Optimality

- › X is a **Pareto improvement** over Y iff someone prefers X to Y , and no one prefers Y to X . So we can rephrase:
Reformulated Pareto If X is a Pareto improvement over Y , then X is morally better than Y .
- › An outcome is **Pareto efficient** or **Pareto optimal** iff there is no outcome that is a Pareto improvement over it.
- › Typically many outcomes will be Pareto efficient; they may not be ‘optimal’ in any ordinary sense of that term:

there are infinitely many possible allocations of property rights, and infinitely many allocations of goods and services, that meet the definition of Pareto optimality. A highly egalitarian allocation can be Pareto optimal. So can any allocation where one person has all the wealth and everyone else is reduced to a bare subsistence. (Quiggan 2019: 147-48)

Acting on Pareto

- › How can we **operationalise** the Pareto principle, and turn it into an action-guiding evaluative tool?
- › The idea is that you can determine preferences empirically (see **lecture 2**), and if you can foresee the **consequences** of actions, you can then tell if the outcome of a course of action is a Pareto improvement on the *status quo*.
- › If it is, you are **permitted** to perform it.
- › There are very few moral obligations on this view – even though we are morally required to make things morally better, the only cases in which the Pareto principle guarantees there is a morally better outcome are those where the *status quo* is not Pareto optimal.
- › Utilitarianism is a **demanding** theory, since it requires that you perform a utility-maximising act; as so far developed, welfare economics demands only that you seek Pareto optimality.
- › **It doesn't even provide guidance on which Pareto optimal outcome you should aim at!**

Consequentialist Welfare Economics

- › Because Pareto optimality is a feature of outcomes, we can see that this theory will be **consequentialist** (lecture 8).
- › In particular, it will be:
 - › **Direct**, because acts are evaluated in terms of their outcomes;
 - › **Non-maximising**, because there is no grounds to rank optimal outcomes against each other;
 - › **Actual**, because an action is permissible iff its actual outcome is Pareto-improving.
 - › **Agent-neutral**, because preference satisfaction is non-perspectival;
- › It is not clear yet whether it is **welfarist**, because it doesn't provide a theory of value: just a **sufficient condition** for 'better than'.
- › The theory is thus at best a **partial** account of the 'problem of policy' (Samuelson 1950: 11).

Optimality and Competition

- › The notion of Pareto optimality is closely connected with **perfectly competitive markets**.
- › The **first fundamental theorem of welfare economics** says that, in a perfectly competitive market, the allocation of goods that results will be Pareto optimal (Reiss 2013: 212; Hausman, McPherson, and Satz 2017: 148).
 - › This can sound impressive - ‘competitive markets are optimal!’
 - › But really it just means that, if you assume there is always an opportunity for a Pareto improvement, and sufficiently many market participants exist to permit free and non-disadvantageous trade to realise that improvement, then the resultant market condition after all trades have completed will be Pareto optimal - this is almost ‘trivial’ (Reiss 2013: 212).
- › The theorem entails that the ‘optimal’ allocation cannot be improved upon without someone losing out; that **doesn’t** say that it cannot be improved upon!
 - › Is a Pareto improvement always a genuine improvement?

Optimality and Competition: role of initial allocations

- › The **second fundamental theorem of welfare economics** says (more or less) that every Pareto optimal outcome is obtainable from **some** initial allocation of property rights in a perfectly competitive market (Reiss 2013: 212; Hausman, McPherson, and Satz 2017: 149).
- › This simply underscores the **ethical incompleteness** of the account.
- › The policy evaluation of the initial allocation is not constrained **at all** by the notion of Pareto optimality – the ‘predistribution’ that determines the optimal market outcome needs to be justified in some other way (Quiggan 2019: ch. 12) – if it is even able to be altered.
- › Pareto optimality is incomplete too when we consider goods that ought not to be allocated by market mechanisms: perhaps ‘body parts or sexual services’ (Hausman, McPherson, and Satz 2017: 150). If there were a perfect market in such goods, it would be immoral but optimal. (This is a major topic of **lecture 11**.)
 - ›› To say this is to simply deny that the Pareto principle ‘better than’ ordering is aligned with the **real** better-than relation.

Cost-Benefit Analysis

Cost-Benefit Analysis

This is a bit of a mess.

Questions unanswered

Economists need to be able to make comparisons between states of affairs where neither alternative is Pareto superior to the other. Pareto improvements are rare. Economic changes usually involve winners and losers, and it is not a matter of social indifference who wins and who loses. Suppose that there is a single consumption good (bread, say) in some fixed quantity and that everybody prefers more rather than less of it. Then every distribution of bread that exhausts the bread supply is Pareto efficient, and none of these distributions is a Pareto improvement over any other. Moreover, R may be Pareto optimal and S may be suboptimal without R being a Pareto improvement over S . Suppose there are only two people, A and B , and 10 loaves of bread to distribute. A Pareto efficient allocation that gives 7 loaves to A and 3 to B is not a Pareto improvement over the (suboptimal) allocation that wastes 2 loaves and gives 4 to both A and B , and the Pareto principle says nothing about which is better. (Hausman, McPherson, and Satz 2017: 152)

Extending the ordering

- › If we want to address these incompleteness worries, we need to extend the ordering to more outcomes.
- › The standard approach is to make use of **cost-benefit analysis**, which **emulates** the way utilitarians model **trade-offs**.
- › The utilitarian splits total welfare from the details of distribution; when trade-offs lead to an improvement in aggregate wellbeing, they are acceptable, regardless of distribution. (There will be potentially many utility maximising allocations.)
- › Welfare economics says: **do the same!** When there is a Pareto improvement, there is an improvement in total wellbeing for the welfare economist.
- › So if we have some way of **grouping outcomes** so they have the same quasi-total wellbeing as a Pareto improvement, differing only in distribution, then welfare economics says those outcomes also are better than the *status quo*, even if they are not Pareto improvements because some trade-off involves some people being less satisfied than currently.

Capacity

- › The proposal is that the relevant notion of ‘quasi-total wellbeing’ is ‘quantity of economic benefits to be distributed – ... their *capacity to satisfy preferences*’ (Hausman, McPherson, and Satz 2017: 159).
- › To avoid interpersonal differences in the values assigned to different goods, it is standard to aggregate these quantities by **types of goods**.
 - › Let’s say that in some family of possible outcomes, there are two goods to be distributed: A and B .
 - › Each outcome is some distribution of some total quantity of these goods.
 - › Two outcomes, X and Y , have the same **capacity**, which I write $X \simeq Y$, iff they have the same total amount of A and the same total amount of B .
 - › We don’t need the idea that any outcome has a well-defined capacity in itself, or how to weight and sum quantities of A and B – only that we can tell when two outcomes have the same capacity.
- › When two outcomes have the same capacity, the total goods to be distributed is constant; only who gets how much differs.
 - › If there are 10 loaves of bread to be distributed, all the optimal allocations which exhaust the bread supply have the same capacity.

Capacity in Application: Cost-Benefit Analysis

- › In practice, this framework assumes that one of the goods is money, so that the different distributions with the same capacity involves allocating goods to one individual and compensating money to the other.

Cost-benefit analysis (CBA) seeks to evaluate socio-economic projects in terms of the net benefit the realization of the project would provide. Net benefit here is understood as the total amount the supporters of a project would be *willing to pay* to realize the project minus the total amount that those who are against the project would require in compensation for agreeing to it. CBA assumes that willingness to pay indicates preference satisfaction. Therefore, the project with the greatest net benefit is regarded as the most efficient at satisfying preferences, and on the preference-satisfaction theory of well-being, the same project is the one that has the greatest capacity to make people better off. (Reiss 2013: 226)

Potential Pareto Improvement

- › To understand this, we need an extended notion of an improvement:

Potential Pareto Improvement If (i) X' is a Pareto improvement over Y , and (ii) $X \simeq X'$, then X is a **potential Pareto improvement** over Y .

- › And we can now tweak the Pareto Principle:

Potential Pareto If X is a potential Pareto improvement over Y , then X is morally better than Y .

Bread Revisited

Table 1: Some possible allocations of loaves of bread

| Individual | Outcome | | | |
|--------------|----------|-----------|-----------|-----------|
| | Y | X_1 | X_2 | X_3 |
| Alice | 4 | 5 | 3 | 0 |
| Bob | 4 | 5 | 7 | 10 |
| Total | 8 | 10 | 10 | 10 |

- › In table 1 X_1 is a Pareto improvement over Y , making everyone better off in bread.
- › The idea is that therefore X_2 and X_3 are also better than Y , even though some people are worse off in terms of their preferences, because the overall capacity is improved.
- › Note that, just as utilitarians don't care about distributions of wellbeing, this proposal doesn't care about distributions of goods – a big difference!

Interdependence of Equity and Efficiency

- › Despite its appeal, this doesn't really provide a coherent extension of the 'better than' ranking.
- › Suppose Alice and Bob are **accidentally egalitarian**: they don't like having more than other people, it makes them feel undesirably uncomfortable.
- › They both prefer X_1 to Y ; they also prefer Y to X_3 .
- › This is a problem:
 - › X_1 is a Pareto improvement over Y ; so X_3 is a potential Pareto improvement over Y ; so X_3 is strictly better than Y .
 - › Y is a Pareto improvement over X_3 ; so Y is strictly better than X_3 .
- › We've ended up with an **inconsistent** ranking of outcomes!
- › Because people may have preferences that are correlated with facts about distributions,
 - there is no general way to separate questions concerning the capacity to satisfy preferences from distributional questions. (Hausman, McPherson, and Satz 2017: 169)

A General Theoretical Picture




Figure 1: Preferences between possible allocations (Samuelson 1950: fig. 6)

(A 'utility possibility function' gives the possible joint utilities for two agents for a given total capacity.)

Compensation and Cost-Benefit Analysis

- › This kind of case can be avoided only if the utility-possibility curves for two ‘capacity classes’ do not cross.
- › Otherwise, while winners and losers out of a capacity shift **could** be compensated, there is no guarantee they **will** be; and why should this assuage the objections of the losers?
 - › Take seriously the welfare economist’s objections to interpersonal welfare comparisons: unless you shift to a Pareto improvement, you seem to be prioritising some people’s preference satisfaction (i.e., wellbeing) over others!
- › It doesn’t even show they **could** be compensated: in real cases, there are **transaction costs** which impede shifts from one distribution to another:
 - we could move people to different points on the utility-possibility function only by an ideally perfect and unattainable system of absolutely lump-sum taxes or subsidies. (Samuelson 1950: 18)
 - › Transaction costs hit the poor more than the rich – outcomes which may improve things for the poor may be blocked because the **feasible** redistributions are not Pareto efficient.

Welfarism

- › From a welfarist perspective, even in the best case the approach doesn't yield reliable policy recommendations:

The inadequacy for actual policy decisions—even in the most idealized, simplified world—of all of the discussed measures of 'real income' can be illustrated by numerous examples. Consider the very best case where we can establish the fact that situation 2 is potentially better than 1 (in the sense of having a uniformly farther-out utility-possibility function). Would a good fairy given the chance to throw a switch from 1 to 2 be able to justify doing so? Upon reflection we must, I am afraid, answer no. Potentialities are not actualities—and unless she can give a justification of her act that will satisfy all reasonably defined social welfare functions, she cannot know whether or not to pull the switch. (Samuelson 1950: 11)

- › The problem: what if an outcome is unacceptable on some **social welfare function that cares about how wellbeing is distributed** (Hausman, McPherson, and Satz 2017: 165)?

Costs and Benefits In Practice

- › The application of the framework of welfare economics is, in practice, only as good as the evidence about preferences that goes into it.
- › There are reasons to suspect that preference is only **imperfectly revealed** in actual consumption, and hence the cost-benefit approach is best applied in only a narrow range of cases:

it would be unreasonable to apply it indiscriminately even as a rule of thumb. A cost-benefit rule will likely be less worrying in cases where the winners and losers are close to equal in ability to pay [low impact of feasibility], where uncertainty about consequences is limited, where people's preferences reflect their judgment concerning which policies best serve their interests, and where most people understand the kinds of reasons recommending different policies. (Hausman, McPherson, and Satz 2017: 168)

- › Nevertheless, given that better information about wellbeing is often no easier to identify, there may be 'no better alternative' (Hausman, McPherson, and Satz 2017: 170), despite the imperfections of cost-benefit analysis in practice.

References

References

- Hausman, Daniel, Michael McPherson, and Debra Satz (2017) *Economic Analysis, Moral Philosophy, and Public Policy*, 3rd edition. Cambridge University Press.
- Quiggan, John (2019) *Economics in Two Lessons*. Princeton University Press.
- Reiss, Julian (2013) *Philosophy of Economics*. Routledge.
- Samuelson, Paul A (1950) 'Evaluation of Real National Income', *Oxford Economic Papers* 2: 1–29. doi:[10.1093/oxfordjournals.oep.a041383](https://doi.org/10.1093/oxfordjournals.oep.a041383).