Indecisiveness in behavioral welfare economics☆

Michael Mandler

Department of Economics, Royal Holloway College, University of London, Egham, Surrey TW20 0EX, United Kingdom

ABSTRACT

When an individual's preferences depend on the time or 'frame' at which decisions are made, the preferences that appear at different frames must be aggregated in order to make social decisions. Suppose we aggregate each individual i's frame-based preferences with a 'behavioral welfare relation' that ranks x above y if, when both x and y are available, i sometimes chooses x and not y and never chooses y and not x. The set of Pareto optima can then be large. In fact a small amount of preference diversity across frames can cause every allocation to be Pareto optimal. More generally, the set of Pareto optima will have the same dimension as the set of allocations. The Pareto criterion then will not be able to discriminate locally among policy options. A small distortion, for example, will call for no policy response.

© 2013 Published by Elsevier B.V.

1. Introduction

In the face of individual behavior that violates the principles of economic rationality, behavioral economists often disaggregate an individual into a set of agents acting at different times or 'frames'. An individual who displays an endowment effect becomes a set of agents, one preference relation for each endowment (Tversky and Kahneman, 1991); a hyperbolic discounter becomes a set of agents, one utility function for each point in time; and so on. But if an individual is viewed as a set of agents, the definition of individual welfare becomes problematic. When the agents who choose at different frames disagree on how to rank outcomes, which agent rules? The natural answer, for economists, is to apply the Pareto criterion. In Bernheim and Rangel's (2007, 2009) welfare theory, an individual is weakly better off with x than y if there is no frame at which the individual chooses y but not x when both are available. Salant and Rubinstein (2008) propose a similar model though geared to the positive task of seeing when an individual's choices can be explained as the outcome of distinct rational agents choosing at separate frames. In a spirit similar to Bernheim and Rangel, Mandler (2004, 2005) argues that a unified view of individual welfare can be preserved if an individual is described by an incomplete preference relation: the individual's preference judgments consist only of those rankings on which the disaggregated agents unanimously agree.

In this paper, I consider whether a welfare economics built on these foundations is sufficiently decisive. The above models all take the view that if the disaggregated versions of an individual i who choose at different frames disagree about how to rank outcomes x and y then there is no unified ranking of x and y for individual i. Accordingly, the 'behavioral welfare relation' $\succsim_i$ that we define for individual i will be incomplete.

For a society of individuals, each with a behavioral welfare relation, the incompleteness of the $\succsim_i$ can cause the set of Pareto optima to be very large, thus crippling the Pareto criterion's capacity to rank allocations. In fact the introduction of only a little diversity of behavior across frames at the individual level can sometimes lead every allocation to be Pareto optimal. More generally, the set of Pareto optima will have the same dimension as the set of allocations; every allocation in the neighborhood of an optimum will then be another optimum. Consequently the Pareto criterion does a poor job of locally discriminating among allocations of goods. For example, if the government selects a Pareto optimum and the model is perturbed slightly, say by the introduction of a small externality, then typically the initial optimum will remain optimal. Even paradigmatic distortions do not call for a policy response.

The expansion of the set of Pareto optima stems from the incompleteness of the behavioral welfare relation $\succsim_i$; incompleteness makes it harder to find Pareto improvements and hence easier to declare an allocation optimal. But not any variety of incompleteness will lead to trouble. The key ingredient is that agents' behavioral welfare relations have multiple supporting prices, which in geometric terms means that the boundary of the set of consumption bundles that are $\succsim_i$-superior to an arbitrary bundle will be kinked. We will argue through a series of examples in Section 4 that this pattern is common in behavioral models when an individual has conflicting preferences at different frames. Even in intertemporal choice, which in Bernheim and Rangel (2009) at first appears to be a case where their welfare theory applies fruitfully, the kinked pattern emerges and there will be a large set of Pareto optima.

To understand the link between multiple supporting prices and large sets of Pareto optima, suppose the individuals in society have behavioral welfare relations whose sets of supporting prices have the maximum possible dimension and that those sets robustly overlap at some Pareto optimum. Then normally any slight change in the allocation of goods will lead to sets of supporting prices that continue to overlap. By a version of the first welfare theorem, this new allocation must also be Pareto optimal. We will also show, though this argument is more difficult, that most Pareto optima do in fact exhibit a robust overlap of supporting prices and that the only exceptions occur on the boundary of the set of optima.

That the incompleteness of behavioral welfare relations leads to some indecisiveness in welfare rankings is to be expected. But there are several surprising points. First the dimension of set of optima will normally be large – the largest dimension possible in typical applications. Second, on top of the dimensional expansion, the size of the set of optima can be large. The introduction of a little diversity across frames of each individual's preferences can even cause every allocation to become optimal, as we show in an example in Section 6. Finally the dimensional expansion of the set of optima is 'generic': almost every optimum is surrounded by a full-dimensional set of other optima.

Bernheim and Rangel (2009), responding to an earlier version of the results in this paper, argue that when the behavioral diversity across frames is small the size of set of optima will be small even when its dimension is large. Our example in Section 6, where a small amount of individual-level diversity leads the entire set of allocations to be optimal, shows that this argument is not correct. Focusing on cases with a low level of diversity across frames is also questionable: the need for a behavioral welfare economics is driven by the fact that deviations from classical choice theory are rampant.

Formally, our analysis of the dimension of the set of Pareto optima can be detached from the fact that each behavioral welfare relation $\succsim_i$ originates from a set of disaggregated agents choosing at different frames. Our results apply to any model with preferences that meet the multiple-supporting-prices assumptions that behavioral welfare relations typically satisfy. In this respect, several of our points follow in the footsteps of Rigotti and Shannon (2005). With the aim of showing that competitive equilibria are indeterminate, Rigotti and Shannon consider the Pareto set in economies with uncertainty where, as in Bewley (2002), preferences are incomplete due to the fact that agents possess sets of probability distributions rather than a single distribution. See also Dana (2004) for a similar but more specific model. Our characterization of the optimal allocations via intersecting sets of supporting price vectors in Section 7 parallels Rigotti and Shannon. See Billot et al. (2000) for a similar construction, and also Bonnisseau and Cornet (1988).

We differ from the above literature in minor and major ways. The minor way is that our model does not use probabilities and hence lacks the linearity assumptions of the Bewley framework.[1] One major deviation is our goal of showing that there is *typically* a large multiplicity of optima (that a large multiplicity obtains in the neighborhood of almost every optimum). To tackle this point, we need to face the technical hurdle that boundary Pareto optima inevitably arise where agents' sets of supporting prices are tangent and do not overlap robustly. In these troublesome cases, a Pareto optimum need not be surrounded by a full-dimensional (open) set of other optima. However a global analysis of the Pareto optima, in Section 8, will show that generically the troublesome cases constitute only a measure zero subset of the Pareto optima. A second major difference arises when we consider behavioral welfare relations with sets of supporting prices that are multidimensional but fall short of full dimensionality (see Appendix B). While the project of showing that competitive equilibria are indeterminate will then break down, the dimension of the set of Pareto optima will still be large.

The expansion of the set of optima discussed in this paper is a characteristic drawback of the Paretian aim of avoiding inter-agent comparisons of welfare. Even without the complications of frame-based selves, Pareto optimality is a problematic guide to policymaking. Since policymakers never know with certainty the preferences of individuals, a Pareto improvement that avoids imposing interpersonal comparisons of utility must be an improvement for each of the preference relations that any individual might potentially have. This test is so difficult to pass that it becomes nearly impossible to find Pareto

---

[1] Rigotti and Shannon point out that their results do not hinge on these linearity assumptions.