Indecisiveness in behavioral welfare economics

Michael Mandler*

Royal Holloway College, University of London

July 2008

Abstract

Suppose an agent is modeled by a 'behavioral welfare relation' that aggregates the different preferences that the agent reveals when choosing at different times or frames: x is better than y if the agent never chooses y when x is available and sometimes chooses x. In many behavioral applications, the options that are ranked superior to an alternative by this relation will be supported by multiple supporting price vectors. As a consequence, in a society of such agents the set of Pareto optimal allocations can be large and even have the same dimension as the set of all allocations. A policymaker will then not be able to use Pareto optimality to discriminate locally among allocations. A small distortion, for example, will call for no policy response.

JEL codes: D11, D60, C60.

Keywords: behavioral welfare economics, incomplete preferences, Pareto optimality.

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

1 Introduction

In the face of individual behavior that violates the principles of economic rationality, behavioral economists have often disaggregated 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. One drawback of this disaggregation is that the definition of welfare becomes problematic. If 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. Bernheim and Rangel (2007, 2008) offer a detailed theory of welfare economics, applicable in principle to any choice behavior, that infers that an individual is weakly better off with x than y if and only if no agent at any frame chooses y over x and at least one agent at one frame does choose x over y. 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 product of rational agents choosing at separate frames. Mandler (2004a, 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 about which the disaggregated agents unanimously agree.

In this paper, we 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 individual i does not have a welfare ranking of x and y; we will then say that individual i's 'behavioral welfare relation' $\succsim_i$ is incomplete. For a society of individuals, each with a behavioral welfare relation, the incompleteness of the $\succsim_i$ can cripple the Pareto criterion; the set of Pareto optimal allocations can be very large and indeed can have the same dimension as the set of allocations. Since the Pareto criterion then locally makes no discriminations, it does a poor job of identifying acceptable allocations of goods: every allocation in the neighborhood of an optimum is another optimum. Moreover, policymakers need not respond to changes in the environment. For example, if the government selects a Pareto optimum and the model is perturbed slightly, say by the addition of a small externality, then typically the initial optimum will remain optimal: even a paradigmatic distortion does not call for a policy response.

The expansion of the set of Pareto optima stems from the incompleteness of the behavioral welfare relation $\succsim_i$, which makes it difficult to find Pareto improvements and hence easier to declare an allocation Pareto optimal. But not any variety of incompleteness will lead to trouble. The key ingredient is that agents' behavioral welfare relations have multiple supporting prices: the boundary of the set of bundles that are $\succsim_i$-superior to an arbitrary bundle must be kinked. We will argue through a series of examples that this pattern is common in behavioral models where the agents who choose at different frames have divergent preferences.

Formally, the characterization we give 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 general equilibrium model with preferences that meet the multiple-supporting-prices assumptions that behavioral welfare relation satisfy. In this regard, several of our points follow in the footsteps of Rigotti and Shannon (2005). They characterize the Pareto set in economies with uncertainty where preferences are incomplete because agents possess sets of probability distributions, as in Bewley (1986) (see also Dana (2004) for a similar but more specific case). Our characterization of the optimal allocations via intersecting sets of supporting price vectors in section 4 parallels this treatment. Billot et al. (2000) use a similar construction. See also Bonnisseau and Cornet (1988).

The large multiplicity of the Pareto set in the neighborhood of specific allocations could be deduced from the Rigotti and Shannon (2005) treatment. Our goal, however, is to show that a large multiplicity obtains in the neighborhood of almost every optimum. We therefore face the technical hurdle that boundary Pareto optima inevitably arise where agents' sets of supporting prices 'just' overlap, i.e., where the interiors of agents' sets of supporting prices do not intersect. In these troublesome cases, a Pareto optimum need not be surrounded by a full-dimensional (open) set of other optima. But a global analysis of the Pareto optima, in section 5, will show that generically the troublesome cases constitute only a measure zero subset of the Pareto optima. Our analysis mostly assumes that agents' sets of supporting prices is full-dimensional, as in Rigotti and Shannon. In section 6, we consider the dimension of the optima that occurs when sets of supporting prices have smaller dimension.

2 Behavioral welfare and multiple supporting prices

We consider an agent who chooses from a variety of sets A, where each A is a subset of some family of conceivable choice alternatives X. Each time the agent selects from some A an 'ancillary' condition or 'frame' f is present that can affect the agent's behavior. Let $c(A, f)$ be the nonempty subset of A that the agent chooses when the frame is f and let $\mathcal{F}$ denote the domain of (A, f) pairs for which the agent's choices can be observed. A prime case of a frame occurs when choice is affected by some alternative in A that the agent views as the status quo. In this case, f denotes the status quo option and each $(A, f) \in \mathcal{F}$ must have $f \in A$.

We have adopted the 'frame' terminology and notation from Salant and Rubinstein (2008). The following definition of behavioral welfare, the main subject of this paper, is a slight adaptation of the analogous concept in Bernheim and Rangel (2008).

Definition 1 *The behavioral welfare relation $\succsim$, a binary relation on X, is defined by*

$$x \succsim y \iff \begin{cases} y \in c(A,f) \Longrightarrow x \in c(A,f) & \textit{for all } (A,f) \in \mathcal{F} \textit{ with } x, y \in A\,, \\ x \in c(A,f) & \textit{for some } (A,f) \in \mathcal{F} \textit{ with } x, y \in A\,. \end{cases}$$

So $x \succsim y$ obtains if, whenever x and y are both available, we never see y chosen without x also being chosen and x is sometimes chosen. The *strict behavioral welfare relation* $\succ$ is defined by $x \succ y \iff (x \succsim y$ and not $y \succsim x)$. So we make the welfare inference $x \succ y$ if the agent sometimes chooses x and not y and never chooses y and not x. Think of each decision at each (A, f) as being made by a separate agent and interpret the observation $(x \in c(A, f),\ y \in A)$ to indicate that the (A, f) individual weakly or strictly prefers x over y according to whether $y \in c(A, f)$ or $y \notin c(A, f)$. Then $x \succ y$ obtains only when all of the agents who reveal a preference between x and y at least weakly prefer x and at least one strictly prefers x: in this sense, $x \succ y$ indicates a Pareto improvement for the individual's frame-based selves.[1]

Even when agents are not classical economic maximizers, it is often straightforward to link an observation of $x \succ y$ to the agent's belief that he or she experiences greater welfare with x than y. Suppose that an agent endowed with a status quo option is offered the chance to switch to some alternative. Then a (A, f) takes the form $(\{x, y\}, x)$, where x is the status quo, acceptance of the offer to switch is indicated by $c(\{x, y\}, x) = \{y\}$ and rejection is $c(\{x, y\}, x) = \{x\}$. Suppose the agent, unlike a classical maximizer, is sometimes is unable to form a preference judgment and is instead governed by inertia, sticking to the status quo. Then, if we observe $\{x\} = c(\{x, y\}, x)$ we would not have grounds to infer the agent weakly or strictly prefers x to y, just that it is not the case that the agent prefers y to x. But when we see the agent actively repudiate the status quo, $\{y\} = c(\{x, y\}, x)$, and if in addition $\{y\} = c(\{x, y\}, y)$, then it is reasonable to infer that the agent does strictly prefer y to x. These common-sense inferences match the $\succsim$ definition of individual welfare: if $c(\{x, y\}, x)$ and $c(\{x, y\}, y)$ agree – say this option is y – then we deduce $y \succ x$ while if $c(\{x, y\}, x) \neq c(\{x, y\}, y)$, then we infer neither $x \succsim y$ nor $y \succsim x$. So in this example $\succsim$ is simply the union of the preference judgments that the agent can be unambiguously observed to make.

Given the Pareto-like nature of $\succsim$, the incompleteness of $\succsim$ that can occur in the above example is typical. In a more general setting where A can be an arbitrary subset of X, all that is needed for neither $x \succsim y$ nor $y \succsim x$ to obtain is that there is a (A, f) with $(x \in c(A, f), y \in A, y \notin c(A, f))$ and a (A', f') with $(y \in c(A', f'), x \in A', x \notin c(A', f'))$.

[1] There are other plausible definitions for a strict behavioral welfare relation, most prominently, to say that $x \succ^* y$ obtains when $x \succsim y$ and, for all $(A, f) \in \mathcal{F}$ with $x, y \in A$, $y \notin c(A, f)$, that is, when the individual sometimes chooses x when y is available and never chooses y when x is available. Bernheim and Rangel (2008) advocate a strict behavioral welfare relation that, under a minor domain restriction on c, coincides with $\succ^*$. Notice that if $x \succ^* y$ then $x \succ y$, but not vice versa. It will therefore be easier, when we consider societies of many individuals, to find Pareto improvements when $\succ$ rather than $\succ^*$ defines strict improvements for individuals. Hence it will be *harder* to declare an allocation to a Pareto optimum when $\succ$ defines strict improvements. This is our sole reason for using $\succ$ rather than $\succ^*$: to give a convincing case that the set of Pareto optima is large, we must when possible bias the playing field to minimize the number of Pareto optima.

For our main point – that $\succsim$ can make so few rankings that Pareto optimality will not discriminate adequately among policy options – simple incompleteness is not enough. A particular form of incompleteness, where multiple price vectors support the bundles that are $\succsim$-superior to any given reference bundle, is required.

We henceforth narrow our focus to agents who choose bundles of L goods: X will be $\mathbb{R}^L$ or $\mathbb{R}^L_+$.

In standard consumer theory, if a preference relation R on $\mathbb{R}^L$ has a concave utility representation then, given an arbitrary bundle x, there is a 'price' vector p such that yRx implies $p \cdot (y - x) \geq 0$ and we then say that 'p supports the bundles that are R-superior to x.' If u is differentiable at x then p must be a multiple of $Du(x)$. Since concave functions on $\mathbb{R}^L$ are differentiable at almost every point in their domain, there will typically be only one p (up to multiplication by a scalar) that supports the R-superior bundles.

The situation is different for the behavioral welfare relation $\succsim$: for most bundles x, not just an isolated few where differentiability fails, there will be multiple nonproportional price vectors that support the bundles that are $\succsim$-superior to x. Moreover, when there is more than one nonproportional supporting price vector there will be a continuum of such vectors: if $p \cdot (y - x) \geq 0$ and $p' \cdot (y - x) \geq 0$ for all y with $y \succsim x$ then, for any $\alpha, \alpha' \geq 0$, $(\alpha p + \alpha' p') \cdot (y - x) \geq 0$ for all y with $y \succsim x$. See Figures 1 through 3 for the characteristic kink in the set of $\succsim$-superior bundles when there are multiple nonproportional supporting price vectors.

We turn to four examples that show why multiple supporting prices appear with the behavioral welfare relation $\succsim$. To ensure that $c(A, f)$ is always well-defined, we assume without further remark that, for each $(A, f) \in \mathcal{F}$, A is finite. And, to avoid any extra supporting price vectors that do not stem from the multiplicity of frames, we assume that any utility function that appears at a single frame is differentiable. Bernheim and Rangel (2008) derive $\succsim$ in Examples 2 and 3 but do not consider the set of supporting prices.

Example 1 (the willingness to accept-willingness to pay disparity) There are two commodities, a good y that an agent finds difficult to value, e.g., environmental quality, and money m. So $X = \mathbb{R}^2_+$. A frame f is any status quo point (y, m), and accordingly we assume that the agent can always stick to the status quo (for any $(A, (y, m)) \in \mathcal{F}$,

$(y, m) \in A$). We also suppose that we can offer the agent any alternative to any status quo ($\mathcal{F}$ has all sets of the form $(\{(y, m), (y', m')\}, (y, m))$). Given a status quo (y, m), suppose the agent will pay no more than $p_{\text{pay}} > 0$ for each additional unit of y and will accept no less than $p_{\text{accept}} > 0$ for each unit of y sacrificed. The 'disparity' is that $p_{\text{pay}} < p_{\text{accept}}$ (see Kahneman et al. (1991) for a survey). A $(y', m') \in c(A, (y, m))$ with $y' \geq y$ must therefore have a payment $m - m'$ that is no greater than $p_{\text{pay}}(y' - y)$. Hence $m - m' \leq p_{\text{pay}}(y' - y)$. And a $(y', m') \in c(A, (y, m))$ with $y' \leq y$ must lead to receipts $m' - m$ that are at least as great as $p_{\text{accept}}(y - y')$, which gives $m' - m \geq p_{\text{accept}}(y - y')$. Therefore

$$\begin{aligned} m' - m &\geq -p_{\text{pay}}(y' - y) \quad \text{if } y' \geq y, \\ m' - m &\geq -p_{\text{accept}}(y' - y) \quad \text{if } y' \leq y. \end{aligned}$$

So, for increases in y, $c(A, (y, m))$ is bounded below by a line with slope $-p_{\text{pay}}$ and, for decreases in y, by a line with slope $-p_{\text{accept}}$. Since $p_{\text{pay}} < p_{\text{accept}}$, the boundary is less steep for increases (see Figure 1). The kink in the boundary in Figure 1 is the key geometric feature that all of our examples share.

To check that multiple supporting prices obtain for $\succsim$, suppose that $(y', m') \succsim (y, m)$ holds. Then $(y', m') \in c(A, (y, m))$ for $A = \{(y, m), (y', m')\}$ and therefore the above inequalities must be satisfied. Rewrite the inequalities as

$$(p_{\text{pay}}, 1) \cdot ((y', m') - (y, m)) \geq 0 \quad \text{if } y' \geq y, \tag{1}$$

$$(p_{\text{accept}}, 1) \cdot ((y', m') - (y, m)) \geq 0 \quad \text{if } y' \leq y. \tag{2}$$

Since (1) continues to hold if we replace p_{pay} by any $p > p_{\text{pay}}$ and (2) continues to hold if we replace p_{accept} by any $p < p_{\text{accept}}$, we conclude that if $(y', m') \succsim (y, m)$ then

$$(p, 1) \cdot ((y', m') - (y, m)) \geq 0 \quad \text{for all } p \in [p_{\text{pay}}, p_{\text{accept}}],$$

Thus any $(p, 1)$ with $p \in [p_{\text{pay}}, p_{\text{accept}}]$ or any multiple of such a $(p, 1)$ supports the bundles that are $\succsim$-superior to (y, m), as in Figure 1. ■

Example 2 (coherent arbitrariness) In the 'coherent arbitrariness' model of Ariely, Loewenstein, and Prelec (2003), an agent has a distinct utility function at each 'anchor' or frame f in some finite set F. The agent can consume a bundle of L goods and the

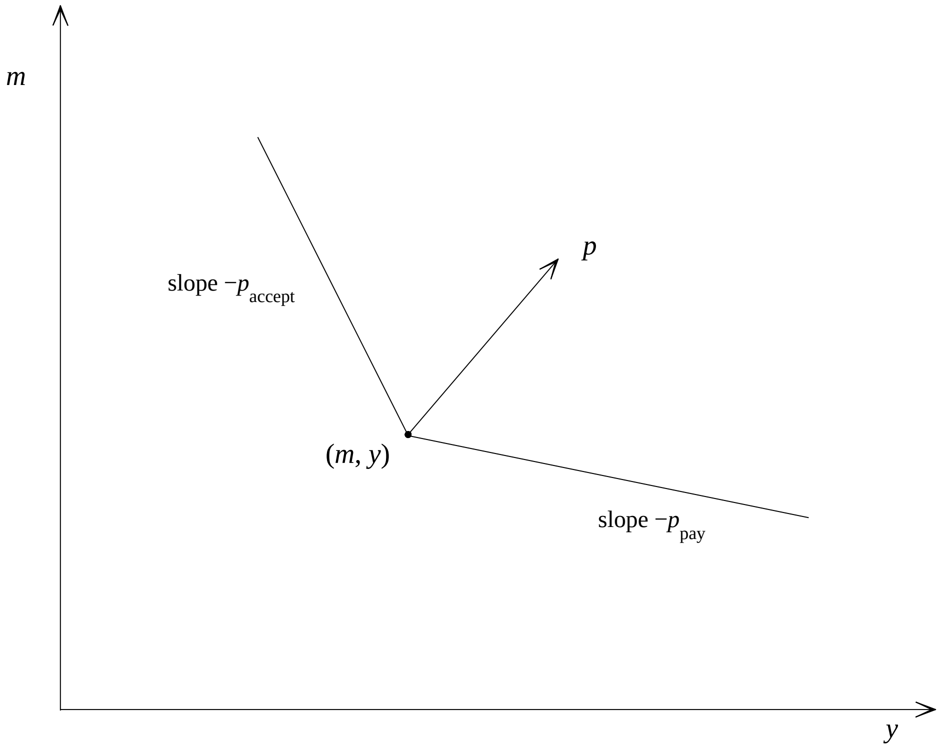

Figure 1: The WTA − WTP disparity

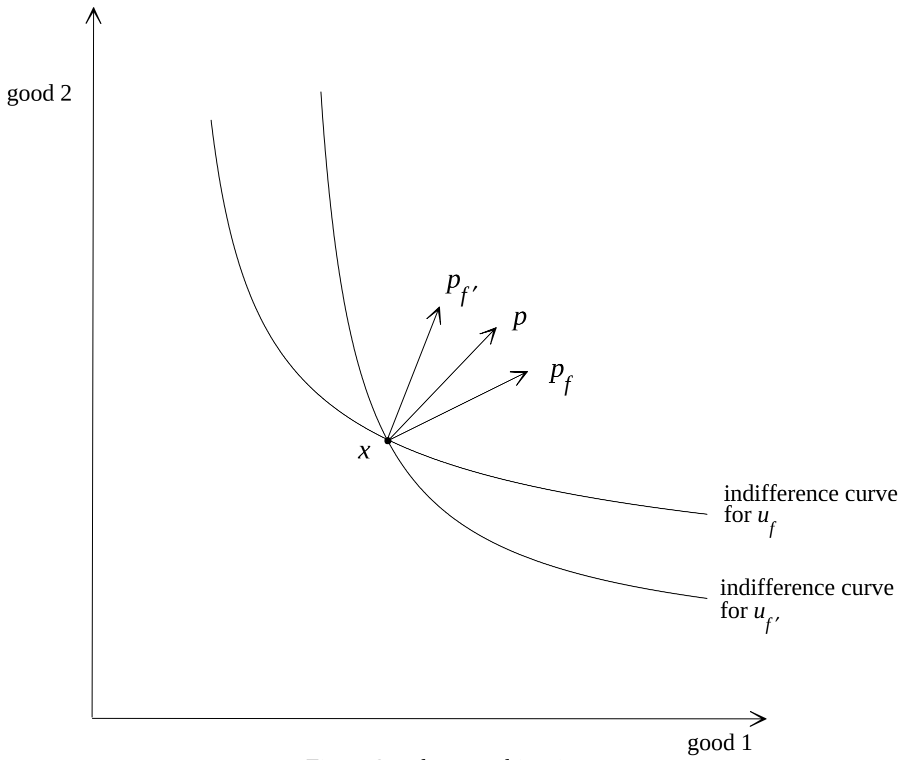

Figure 2: coherent arbitrariness

agent's utility at f is u_f. So we set $c(A,f) = \{y \in A : u_f(y) \geq u_f(x) \text{ for all } x \in A\}$. If $\mathcal{F}$ contains all (A,f) where A has two elements, then, as in Bernheim and Rangel (2008), $\succsim$ is the Pareto relation for the framed utilities: $y \succsim x$ if and only if $u_f(y) \geq u_f(x)$ for all $f \in F$.

Suppose each u_f is quasiconcave. Then, fixing some x, there will be a $p_f \in \mathbb{R}^L$ such that if $u_f(z) \geq u_f(x)$ then $p_f \cdot (z - x) \geq 0$. Thus $z \succsim x$ implies $p_f \cdot (z-x) \geq 0$ for all f. So if p is any positive linear combination of the p_f, $p = \sum_{f \in F} \alpha_f p_f$ where each $\alpha_f \geq 0$, then p supports the $\succsim$-superior bundles, i.e., $p \cdot (z-x) \geq 0$. As long as at least two of the p_f are linearly independent, the multiplicity of supporting prices will not consist of mere rescalings. See Figure 2 for a $L = 2$, $\#F = 2$ example. ■

Example 3 (hyperbolic discounting) An agent consumes a sequence $x = (x_1, ..., x_T)$ from date 1 to $T \geq 3$ with one good per date: $X = \mathbb{R}^T_+$. A frame is a date t at which choice is made. At each t, consumption prior to t has already occurred and so the agent can choose only from sets that have sequences that specify the same consumption from 1 to $t-1$. So we admit (A,t) into $\mathcal{F}$ if and only if t is a date between 1 and T and A is any finite set of bundles in $\mathbb{R}^T_+$ that share the same values prior to t (if $x, y \in A$ then $x_i = y_i$ for $i < t$). Hyperbolic discounting (see Laibson (1997), O'Donoghue and Rabin (1999)) means that there are $\beta, \delta \in (0,1)$ and a differentiable, concave function u such that the date t agent maximizes

$$U_t(x) = u(x_t) + \beta \sum_{i=t+1}^{T} \delta^{i-t} u(x_i).$$

So $c(A,t)$ equals the set of x in A that maximize U_t.

The agents that appear at $t > 1$ reveal preferences only on those subspaces of $\mathbb{R}^T_+$ where the consumption of the goods that appear before t is fixed. So, while the bundles that are U_1-superior to an arbitrary $x \in \mathbb{R}^T_+$ will form a full T-dimensional set, the set of bundles that are U_t-superior to x, where $t > 1$, will have dimension less than T. Consequently, the set of *prices* that support the bundles that are $\succsim$-superior to x will not be restricted by the choices of the agents at dates $2, ..., T$.

The situation becomes more interesting once at least one time period has passed: if $t - 1 > 0$ dates have passed, with $(\overline{x}_1, ..., \overline{x}_{t-1})$ consumed, then the choices of the agents at dates $1, ..., t$ will reveal complete preferences over the remaining $T - (t-1)$ goods that have yet to be consumed. Let bundles now be points in $\mathbb{R}_+^{T-(t-1)}$ and reset the domain of each $U_i, i \leq t$, to be $\mathbb{R}_+^{T-(t-1)}$ by fixing the first $t-1$ coordinates of any bundle to equal $(\overline{x}_1, ..., \overline{x}_{t-1})$. If $y \succsim x$ obtains then we must have $U_i(y) \geq U_i(x)$, $i = 1, ..., t$.[2] Just as in Example 2, for any $x \in \mathbb{R}_+^{T-(t-1)}$ and any $i = 1, ..., t$, there is a $p_i \in \mathbb{R}^{T-(t-1)}$ such that $p_i \cdot (z - x) \geq 0$ for all z with $U_i(z) \geq U_i(x)$. And, again as in Example 2, any p equal to a positive linear combination of the p_i, $p = \sum_{i=1}^{t} \alpha_i p_i$, will support the $\succsim$-superior bundles. While the vectors $p_1, ..., p_{t-1}$ are collinear, p_t will be linearly independent of these vectors: for any agent $j \leq t-1$ and agent t, $DU_j(x)$ is a multiple of $(Du(x_t), \delta Du(x_{t+1}), ..., \delta^{T-t} Du(x_T))$ and $DU_t(x)$ is a multiple of $(Du(x_t), \beta\delta Du(x_{t+1}), ..., \beta\delta^{T-t} Du(x_T))$. So the multiplicity of supporting prices is not a mere rescaling. Bernheim and Rangel (2008, Theorem 11) is a similar result. ■

Example 4 (abstract status quo bias) An agent has a quasiconcave utility function $u : \mathbb{R}^L \longrightarrow \mathbb{R}$ and the frame f is a status quo bundle x. As in Example 1, the agent is always allowed to stick to the status quo (for any $(A, x) \in \mathcal{F}$, $x \in A$) and we can offer the agent any alternative to any status quo ($\mathcal{F}$ contains all sets of the form $(\{x, y\}, x)$). Following Salant and Rubinstein (2008), status quo bias is modeled by the rule: if there is an option in A that gives the agent a utility premium $\beta(x) > 0$ over $u(x)$ then the agent picks the utility-maximizing bundle in A and otherwise the agent sticks to x. So

$$c(A, x) = \begin{cases} \arg\max u(y) \text{ s.t. } y \in A & \text{if } \exists y \in A \text{ with } u(y) > u(x) + \beta(x) \\ x & \text{otherwise} \end{cases}.$$

Now suppose $u(y) > u(x) + \beta(x)$. Then $y \in c(\{x, y\}, x)$ and if $x, y \in A$, we cannot have $x \in c(A, z)$ for any $z \in A$. So $u(y) > u(x) + \beta(x)$ implies $y \succsim x$. Conversely, if $y \succsim x$ and $y \neq x$ then $y \in c(\{x, y\}, x)$ and hence $u(y) > u(x) + \beta(x)$. Given that u is quasiconcave, it follows that for every x there is a p_x such that $p_x \cdot (y - x) \geq 0$ for all y with $u(y) \geq u(x)$. If u is strictly concave, it is easy to use the last two sentences to show that, for any p sufficiently near p_x, $p \cdot (y - x) \geq 0$ for all y with $y \succsim x$ (see Figure 3). We thus have a multiplicity of supporting prices. ■

[2] The agents at dates t and later may impose further restrictions if x and y are equal at some set of coordinates that begin at t, but the t inequalities we have written must always hold.

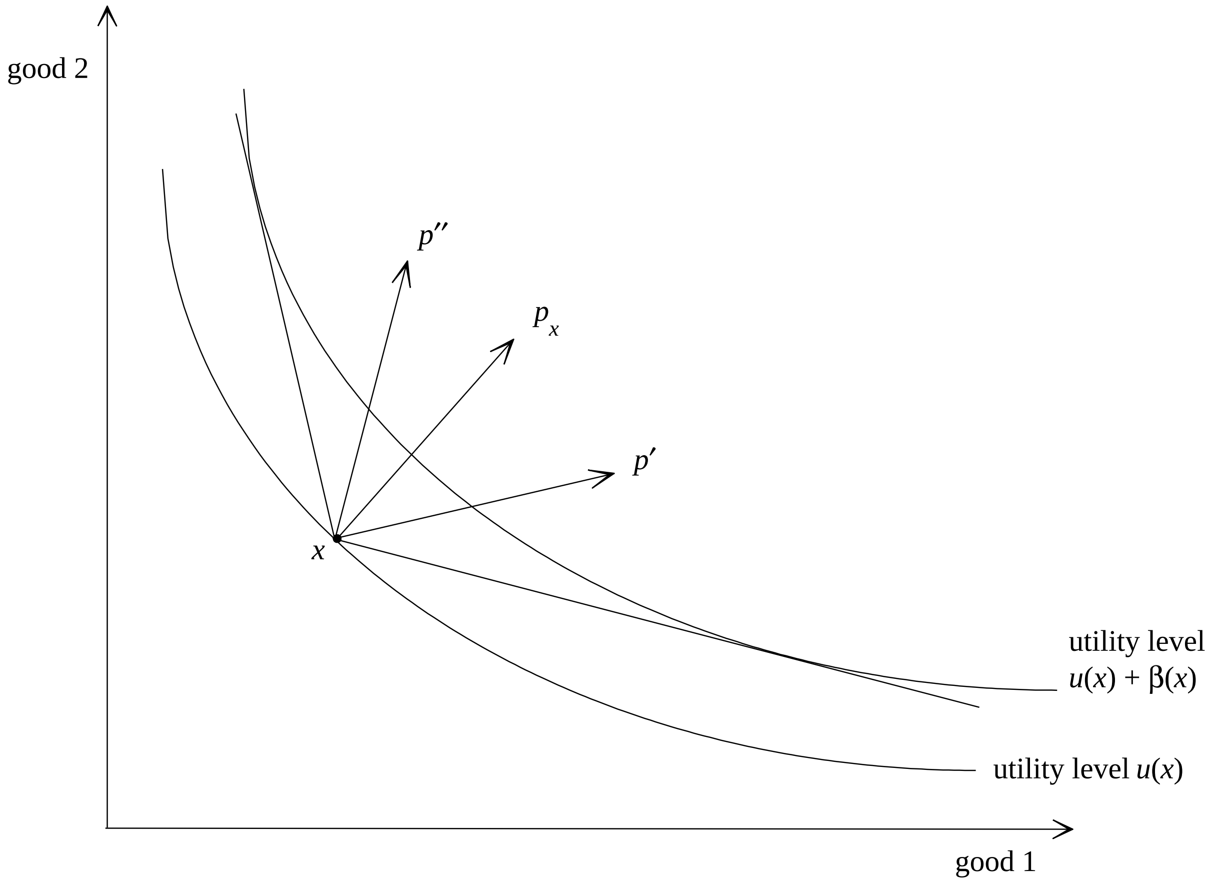

Figure 3: abstract status quo bias.

Any $p = \alpha' p' + \alpha'' p''$ with α', $\alpha'' \geq 0$ supports all y with $y \succsim x$.

Each of the above examples displays a multiplicity of supporting prices that goes beyond the simple rescaling that occurs in standard consumer theory. Since rescalings have no bearing on the set of Pareto optimal allocations, we eliminate them by now requiring any supporting p to lie in the unit circle $S \equiv \{p \in \mathbb{R}^L : \sum_{k=1}^{L} p_k^2 = 1\}$, where L is the number of goods.

The *set of supporting prices at* x for the relation $\succsim$ is given by

$$N(x) = \{p \in S : y \succsim x \Longrightarrow p \cdot (y - x) \geq 0\}.^3$$

We define the *dimension of* $N(x)$ to be the maximum number of linearly independent vectors in $N(x)$ minus 1: any convex combination of these vectors, when rescaled to lie in S, will be an element of $N(x)$. Since S is a $L - 1$ dimensional set, the dimension of $N(x)$ can range from 0 to $L - 1$.

We will concentrate on sets of supporting prices $N(x)$ that, at any x, have the maximum dimension $L-1$. Then, on any two-dimensional plane through x, the set $\{y : y \succsim x\}$ will display a kink at x. The maximum dimension $L - 1$ obtains in Examples 1 and 4. Furthermore, if we extend Example 1 to more than two goods, then $N(x)$ will have dimension $L - 1$ if a WTA-WTP disparity holds for each of the $L - 1$ nonmoney goods and $\succsim$ is convex. For dimension $L-1$ to obtain in Example 2, there must be at least L frames to ensure that there are L linearly independent p_f vectors (and, flukes aside, if there are at least L frames then the dimension will be $L - 1$). Finally, in Example 3 and assuming that $t - 1 > 0$ time periods have passed, the dimension of $N(x)$ is 1. As indicated in the example, the collinearity of the vectors $DU_1(x), ...DU_{t-1}(x)$ constrains the multiplicity of supporting prices.

The moral of the examples is that are cases when it is plausible for the dimension of supporting prices to be the maximum $L-1$, there are cases where a variety of dimensions are plausible, and finally there are cases where the maximum dimension definitely will not obtain. For most of this paper, however, we assume that $N(x)$ has maximal dimension, for the simple reason that the lower dimensional cases introduce no new mathematical or conceptual issues. As we discuss in section 6, if the dimension of $N(x)$ drops from

[3] Each $N(x)$ is closed and the set of unnormalized supporting prices, $\{p \in \mathbb{R}^L : y \succsim x \Longrightarrow p \cdot (y - x) \geq 0\}$, is convex as well as closed.

$L-1$ then the dimension of the Pareto set will drop correspondingly, but the logic of the connection between the dimension of $N(x)$ and the dimension of the Pareto set remains similar.

Definition 2 *The behavioral welfare relation* $\succsim$ *has* maximal supporting prices in a neighborhood of x *if, for every* y *in some open neighborhood of* x*,* $N(y)$ *has dimension* $L-1$*. If* $\succsim$ *has maximal supporting prices in a neighborhood of each* $x \in \mathbb{R}^L_+$ *then* $\succsim$ *has* maximal supporting prices.

A binary relation R that has a set of supporting prices of dimension $L-1$ at a single x just indicates a kink in the boundary of $\{y : yRx\}$, which can occur with a standard preference relation. Indeed, if R is a complete, transitive, and convex preference relation, multiple supporting prices *must* be a rare event: since $\{y : yRx\}$ is convex, a line along this set's lower boundary must be differentiable almost everywhere. The distinctive feature of the behavioral welfare relation $\succsim$, in contrast, is that maximal supporting prices can occur on open sets: the incompleteness of the behavioral welfare relation $\succsim$ allows the multiplicity of supporting prices to be systematic.

3 Further properties of behavioral welfare

Beyond multiple supporting prices, the examples in the previous section show that further assumptions on $\succsim$ can be well-founded. For our purposes, the convexity of $\succsim$ is the most important. Since we cannot assume that $\succsim$ is complete, transitive, or continuous we use a more detailed version of convexity (which, in standard consumer theory, is implied by the traditional convexity assumption).

Definition 3 *A binary relation* $\succsim$ *on a convex set* X *satisifies the* convexity condition *if, for all* $x \in X$*,*

$$x' \succsim x,\ x'' \succsim x,\ \alpha \in (0,1) \implies \alpha x' + (1-\alpha)x'' \succsim x,$$
$$x' \succsim x,\ x'' \succ x,\ \alpha \in (0,1) \implies \alpha x' + (1-\alpha)x'' \succ x.$$

In Example 2, $\succsim$ will satisfy the convexity condition if each u_f is quasiconcave and continuous and, in Example 3, $\succsim$ will satisfy the convexity condition if u is concave.

In these examples, the convexity condition is no more or less implausible than convexity traditionally is. In Example 4, the convexity condition fails (but see our further discussion of this example in section 4).

Definition 4 *The relation* $\succsim$ *is* nonsatiated *if, for all* $x \in X$*, there exists* $y \in X$ *with* $y \succ x$*.*

Examples 2, 3, and 4 will lead to $\succsim$'s that are nonsatiated if the u_f in Example 2 and u in Examples 3 and 4 are increasing.

On other hand, some conditions that are the norm in classical preference theory can fail for $\succsim$ in Examples 1-4. The major technical headache is the failure of continuity. In Example 2, for example, suppose that $u_f(x) > u_f(y)$ holds for all f except for some f' where $u_{f'}(x) = u_{f'}(y)$ holds. Then $x \succ y$ but the set $\{z : z \succ y\}$ need not be open even if each u_f is continuous (consider a sequence of points $x^n \longrightarrow x$ where $u_{f'}(x^n) < u_{f'}(y)$).

We have so far considered specific examples that impose structure on $\succsim$. Does the fact that $\succsim$ is derived from a set of $c(A, f)$ observations entail restrictions on $\succsim$ that we must therefore assume are satisfied? Suppose that for each one- or two-element $A \subset X$ there is a f such that $(A, f) \in \mathcal{F}$. The inclusion of the singleton A's implies that $\succsim$ must be reflexive and so we assume (without further remark) that any behavioral welfare relation $\succsim$ in this paper is reflexive. But there are no further restrictions on $\succsim$. For any reflexive relation R on X, set $x \in c(\{x, y\}, f)$ if and only if xRy and $c(A, f) = A$ for all A with $\#A \geq 3$. Then $\succsim = R$.

4 Pareto optimality

Suppose now there is a society of agents $\mathcal{I} = \{1, ..., I\}$. When a notation from the previous sections carries a i subscript, it refers to the same type of point or set but for the particular individual i. For the remainder of the paper, we fix the set of alternatives $X_i = \mathbb{R}^L_+$ for each $i \in \mathcal{I}$. Each agent $i \in \mathcal{I}$ makes a set of choices $c(A, f)$ for (A, f) drawn from some $\mathcal{F}$, generating a behavioral welfare relation $\succsim_i$.

The society has an endowment of the L goods, $e \in \mathbb{R}^L_{++}$. An *allocation* $x = (x_1, ..., x_I)$ is a point in the $L(I-1)$-dimensional set of feasible allocations $Y = \{x \in \mathbb{R}^{LI}_+ : \sum_{i \in \mathcal{I}} x_i =$

$e\}$. Henceforth the 'boundary' or 'interior' of a set of allocations or an 'open set' of allocations are defined relative to Y.

For $x, y \in \mathbb{R}_+^{LI}$, y *Pareto dominates* x if $y_i \succsim x_i$ for all $i \in \mathcal{I}$ and $y_i \succ x_i$ for some $i \in \mathcal{I}$, and x is a *Pareto optimum* if x is an allocation and there does not exist an allocation y that Pareto dominates x.[4]

We will see that if the $\succsim_i$ satisfy the convexity condition and have maximal supporting prices then typically almost every Pareto optimum resides in an open set of Pareto optima. The initial steps of our argument are based on the classical welfare theorems. First, we show that if x is a Pareto optimum then there is a p that can support the $\succsim_i$-superior bundles for each $i \in \mathcal{I}$; in the language of the second welfare theorem, x is a 'quasiequilibrium.' Then we assume that maximal supporting prices holds and in addition that there is a supporting price vector p that is in the *interior* of each set of supporting prices $N_i(x_i)$. If the $N_i(x_i)$ sets vary continuously as a function of x_i then the same p can continue to serve as a supporting price vector for each i if we move from x to any nearby allocation x'. So these x' allocations are also quasiequilibria. A version of the first welfare theorem then implies that since x' is a quasiequilibrium it must also be a Pareto optimum. In the next section, we justify our focus on the case where p is interior to each $N_i(x_i)$: almost all of the Pareto optima have this feature.

While we rely on welfare theorem arguments, we unfortunately cannot use off-the-shelf versions: the relation $\succsim_i$ can fail to be complete or continuous in even the most-behaved example, and, in principle at least, can fail to be transitive.[5]

An allocation x is a *quasiequilibrium* if there is a p with $p \in N_i(x_i)$ for all $i \in \mathcal{I}$.

Proposition 1 *If each $\succsim_i$ satisfies the convexity condition and at least one $\succsim_i$ is nonsatiated, then any Pareto optimum is a quasiequilibrium.*

Proposition 1, which is a version of the second welfare theorem, makes do with such weak assumptions – no continuity, completeness, or transitivity – because optima are

[4] For results on the existence of Pareto optima in incomplete-preferences environments, see Mandler (2004b); some continuity assumptions are needed. See also Proposition 3.

[5] If we strengthen our nonsatiation assumptions, we could adapt Fon and Otani (1979) to prove Proposition 1.

required only to be quasiequilibria, not full equilibria, and because of our strengthened version of convexity. We have not yet assumed that the $\succsim_i$ satisfy maximal supporting prices.

Let int$N_i(x_i)$ denote the interior of $N_i(x_i)$ relative to S (the set of prices normalized to length 1). Since S has dimension $L-1$, int$N_i(x_i)$ must also have dimension $L-1$ if it is nonempty.

Definition 5 *The relation $\succsim_i$ has* continuous supporting prices *at x_i if $x_i^n \longrightarrow x_i$ and $p \in \text{int}N_i(x_i)$ imply there is a n' such that, for $n > n'$, $p \in \text{int}N_i(x_i^n)$.*

In words, continuous supporting prices holds when a p in int$N_i(x_i)$ cannot jump outside of the set of supporting prices at all bundles arbitrarily near x_i.

Definition 6 *The allocation x is a* robust Pareto optimum *if (1) there is a p such that $p \in \text{int}N_i(x_i)$ for each $i \in \mathcal{I}$, and (2) each $\succsim_i$ has continuous supporting prices at x_i.*

Requiring that a Pareto optimum x is robust is stronger than assuming each $\succsim_i$ has maximal supporting prices in a neighborhood of x_i: there must be among the p's given by Proposition 1 a p that is in the interior of each $N_i(x_i)$.

Proposition 2 *Any robust Pareto optimum is contained in an open and hence $L(I-1)$-dimensional set of Pareto optima.*

The proof of Proposition 2 is straightforward: given a robust Pareto optimum x, both x and any nearby allocation x' can be supported by the same price vector, and hence, by applying a version of the first welfare theorem, x' will also be a Pareto optimum. In the next section, we show that most Pareto optima are robust.

Notice the contrast between Proposition 2 and the size of the Pareto set in the standard general equilibrium model: in an economy of I agents with complete, transitive, strictly convex, and monotone preferences the set of Pareto optima has dimension $I-1$ (Arrow and Hahn (1971)).

Example 5 (abstract status quo bias revisited) Example 4 fails to satisfy our convexity condition and so our arguments in this section and the next do not apply. For this example, however, the conclusion that the set of Pareto optima contains a set of maximal dimension is very simple. Let each $i \in I$ in an economy satisfy abstract status quo bias defined via continuous u_i and β_i functions, leading to the behavioral welfare relation $\succsim_i$. Now suppose x is a Pareto optimum for the economy that consists of the same set of agents $\mathcal{I}$ but where each i has the complete preference relation represented by u_i (so i always chooses a u_i-maximizing bundle without requiring that a change from a status quo x_i delivers a $\beta_i(x_i)$ utility premium). Since a y with $y_i \succsim_i x_i$ for each $i \in I$ has $u_i(y_i) > u_i(x_i)$ if $y_i \neq x_i$, any y that Pareto dominates x in the original economy with abstract status quo bias must also Pareto dominate x in the complete-preferences economy. Such a y must therefore be infeasible and hence x must be a Pareto optimum for the abstract-status-quo-bias economy. Moreover, if $x^n \longrightarrow x$ then for all n sufficiently large x^n must be a Pareto optimum for the abstract-status-quo-bias economy. For if there were a sequence y^n where each y^n Pareto dominates x^n then any accumulation point $\overline{y}$ of y^n must Pareto dominate x in the complete-preferences economy, contradicting the optimality of x.[6] ■

5 The robust optima are generic

We can characterize the robust and nonrobust optima and visualize which case is more likely using the concept of transversal intersection. Suppose that $\succsim_i$ displays maximal supporting prices and is smooth in the sense that the boundary of $N_i(x_i)$ is a smooth $(L-2)$-dimensional surface. If A and B are two subsets of S (the set of price vectors with length 1), they *intersect transversally*, which we write $A \pitchfork B$, if the linear subspaces in $\mathbb{R}^L$ that best approximate A and B at any point of common intersection y together span the linear subspace in $\mathbb{R}^L$ that best approximates S at y.

Consider the simplest case where the economy consists of two agents i and j. If x is a Pareto optimum then $N_i(x_i)$ and $N_j(x_j)$ have a common price vector. If in addition $N_i(x_i) \pitchfork N_j(x_j)$, then x is a robust optimum. To see why, observe that an optimal x will fail to be robust if and only if $N_i(x_i) \cap N_j(x_j)$ consists only of price vectors that

[6] For each agent i, $u_i(y_i^n) \geq u_i(x_i^n)$, and taking a subsequence if necessary there must be an agent j for which $u_j(y_j^n) \geq u_j(x_j^n) + \beta_j(x_j^n)$; hence in the limit $u_i(\overline{y}_i) \geq u_i(x_i)$ for all i and $u_j(\overline{y}_j) \geq u_j(x_j) + \beta_j(x_j)$.

are on the boundaries of both $N_i(x_i)$ and $N_j(x_j)$: in this case, $N_i(x_i)$ and $N_j(x_j)$ are tangent and hence $N_i(x_i) \pitchfork N_j(x_j)$ does not obtain (see Figure 4).[7] Now if we perturb i's and j's sets of supporting prices $N_i(x_i)$ and $N_i(x_j)$ then at any given allocation x a failure of transversal intersection will be an exceptional event. But failures of transversal intersection at *some* x can be unavoidable. As x moves along some path in Y, $N_i(x_i)$ and $N_j(x_j)$ can switch from being disjoint to intersecting transversally, with nontransversal intersection necessarily occurring at some transition point. Since the qualitative fact of the switch from $N_i(x_i)$ and $N_j(x_j)$ being disjoint to intersecting cannot be perturbed away, nontransversal intersection at some allocation will be unavoidable.

But although the Pareto optima where $N_i(x_i)$ and $N_j(x_j)$ are tangent cannot be dismissed, they are unusual in the class of Pareto optima. Consider the set of pairs of allocations and supporting prices for each individual i, $M_i \equiv \{(x,p) \in \mathbb{R}_+^{LI} \times S : p \in N_i(x_i)\}$. We will show (in the proof of Proposition 3) that generically the sets M_i and M_j intersect transversally, $M_i \pitchfork M_j$. Then, although nonrobust optima are not at all pathological, for any nonrobust optimum x there will always be a nearby optimum x' such that $N_i(x_i')$ and $N_j(x_j')$ do intersect transversally. If $N_i(x_i') \pitchfork N_j(x_j')$ failed to hold for an open set of allocations containing x, then the best linear approximations of M_i and M_j at (x,p) would in their p components span only $L-2$ of the $L-1$ dimensions in S, thus contradicting $M_i \pitchfork M_j$. It follows that nonrobust optima appear only on the boundary of the set of robust optima. In fact, we will be able to show that the nonrobust optima have measure 0.

To make the above arguments rigorous, we require the M_i to be manifolds. We say that the relation $\succsim_i$ is *smooth* if M_i is a smooth (C^1) manifold with boundary. To formalize what it means for a model to be typical or generic, we specify a parameter space of economies.

Definition 7 *A* smooth maximal-supporting-prices economy *is an endowment $e \in \mathbb{R}_{++}^L$ and a profile of behavioral welfare relations $(\succsim_1, ..., \succsim_I)$ such that (1) each $\succsim_i$ satisfies the*

[7] A failure of transversality can occur only due to a p that is on boundary of both $N_i(x_i)$ and $N_j(x_j)$. If $p \in N_i(x_i) \cap N_j(x_j)$ and p is in the interior (relative to S) of $N_k(x_k)$ for either $k = i$ or $k = j$ then the linear subspace that best approximates $N_k(x_k)$ at p by itself locally spans all of S at p.

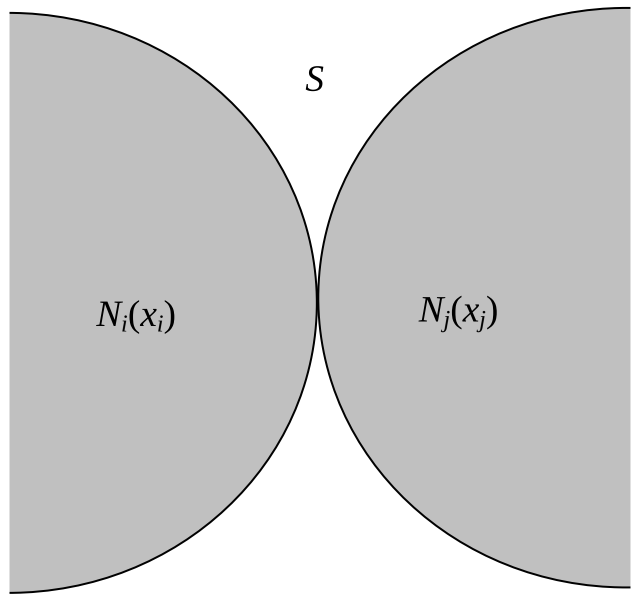

Figure 4: Nontransversal intersection of $N_i(x_i)$ and $N_j(x_j)$

convexity condition, is nonsatiated and smooth, and has maximal supporting prices (2) there exists a Pareto optimum.

Condition (2) avoids some trivialities. Among the weak conditions that ensure the presence of a Pareto optimum are strict monotonicity for one of the economy's I agents (give that agent all of e).

Let M_i be the manifold for $\succsim_i$. A sequence of behavioral welfare relations $\succsim_i^n$ that meet the conditions given in Definition 7 converges to $\succsim_i$ if there is a sequence of C^1 maps $f^n : M_i \longrightarrow \mathbb{R}_+^{LI} \times S$ such that $f^n(M_i) = M_i^n$ (the M_i manifold for $\succsim_i^n$) and f^n converges C^1 uniformly on compacta to the inclusion map of M_i. This definition of convergence then defines open and dense sets of behavioral welfare relations that satisfy (1) in Definition 7 and, using the product topology, of smooth maximal-supporting-prices economies. The convergence of $\succsim_i^n$ requires only that the sets of supporting prices at every bundle converge, not that the behavioral welfare relations themselves converge. We could require this sense of convergence as well, but it is not necessary for our purposes.

Proposition 3 *For an open and dense set of smooth maximal-supporting-prices economies, the robust Pareto optima form a nonempty open set PO_r and thus have positive measure. The remaining Pareto optima are contained in the boundary of PO_r and have measure zero.*

It is noteworthy that the tools of differential topology prove so useful in modeling economies when agents have multiple supporting prices. Although the sets of bundles that are $\succsim_i$-superior to the x_i are inherently kinked, smooth techniques can nevertheless be applied to the prices that support these bundles. See Mas-Colell (1985) for precedents in the theory of production that suggested the present approach.

6 Partial multiplicity of supporting prices

So far we have considered behavioral welfare relations that display *maximal* supporting prices: given some x_i, the bundles that are $\succsim_i$-superior to x_i in any two-dimensional plane through x_i will display a kink at x_i. What of the cases, mentioned in section 2, where the multiplicity of supporting prices concerns only a subset of goods? We show here what it means for there to be multiple supporting prices for only a subset of goods and calculate the dimension of the set of Pareto optima that results. To ensure that the only source of extra dimensions of optima are the goods with multiple supporting prices, we will suppose that the $\succsim_i$, when restricted to the goods that do not display multiple supporting prices, are completely orthodox.

Let p_{-k} denote $(p_1, ..., p_{k-1}, p_{k+1}, ..., p_L)$, and, for $p \in S$, let $S_k(p)$ denote $\{q \in S : \frac{1}{\|q_{-k}\|}q_{-k} = \frac{1}{\|p_{-k}\|}p_{-k}\}$. So a price vector is in $S_k(p)$, if, relative to p, only the relative price of good k is allowed to vary. There is a multiplicity of supporting prices with respect to good k if we can vary the prices that support a bundle x_i within S_k.

Definition 8 *The behavioral welfare relation* $\succsim_i$ *displays* multiple supporting prices for good k at the allocation x *if there is a* $p \in S$ *such that* $\text{int}_{S_k(p)}(N_i(x_i) \cap S_k(p)) \neq \varnothing$. *A Pareto optimum* x *is* robust *if (1) there is a* p *such that, for each agent* i, $p \in \text{int}_{S_k(p)}(N_i(x_i) \cap S_k(p))$ *when* $\succsim_i$ *displays multiple supporting prices for good* k *at* x, *and (2) for each agent* i, *there is an open* $O \subset \mathbb{R}^{LI}_+$ *that contains* x *such that* $\{(x,p) \in \mathbb{R}^{LI}_+ \times S : p \in N_i(x_i)\} \cap (O \times S)$ *is a manifold with boundary.*

Condition (2) is comparable to the continuous supporting prices assumption in section 4.

We have not and do not assume that each individual displays multiple supporting prices for the same set of goods.

Given an allocation z, let $\mathbb{R}^{nmp}_i(z)$ denote the bundles such that only the goods for which $\succsim_i$ does *not* display multiple supporting prices at z are allowed to vary from z_i: $\mathbb{R}^{nmp}_i(z) = \{x_i \in \mathbb{R}^L_+ : x_i(k) = z_i(k)$ for all k for which $\succsim_i$ displays multiple supporting prices at $z\}$. We define i's *conditional preferences* $\succsim_i(z)$ on $\mathbb{R}^{nmp}_i(z)$ by $x_i \succsim_i(z)\ y_i$ if and only if ($x_i \succsim_i y_i$ and $x_i, y_i \in \mathbb{R}^{nmp}_i(z)$). Let $\mathbb{R}^{nmp}(z) = \mathbb{R}^{nmp}_1(z) \times \cdots \times \mathbb{R}^{nmp}_I(z)$. Letting $S^{nmp}(z)$ denote $\{p \in S : p(k) \neq 0 \implies$ for some $j \in \mathcal{I}$, $\succsim_j$ does not display multiple supporting prices for k at $z\}$, each $\succsim_i(z)$ and $x_i \in \mathbb{R}^{nmp}_i(z)$ defines a set of supporting prices $N^z_i(x_i) = \{p \in S^{nmp}(z) : y_i \succsim_i(z)\ x_i \implies p \cdot (y_i - x_i) \geq 0\}$.

Suppose each i does not display multiple supporting prices for all goods at some allocation z. Then, if each i's conditional preferences are sufficiently well-behaved, and we constrain each i's consumption of the goods for which $\succsim_i$ does display multiple supporting prices to equal the levels specified by z, the allocations that are Pareto optimal given these constraints will normally form a set of dimension $I - 1$.

Definition 9 *An allocation z has a* well-behaved conditional economy *if (1) the set of conditional optima, $PO^{nmp}(z) \equiv \{x \in \mathbb{R}^{nmp}(z) : x$ is a Pareto optimum for the economy $(e, \succsim_i (z)_{i \in \mathcal{I}})\}$, is a $I - 1$ dimensional manifold, (2) for each $x \in PO^{nmp}(z)$ the set of supporting prices $\bigcap_{i \in \mathcal{I}} N_i^z(x_i)$ consists of one price vector, (3) if $z^n \longrightarrow z$, $x^n \longrightarrow x$, $x^n \in PO^{nmp}(z^n)$, $p^n \longrightarrow p$, and $p^n \in \bigcap_{i \in \mathcal{I}} N_i^{z^n}(x_i^n)$, then $x \in PO^{nmp}(z)$ and $p \in \bigcap_{i \in \mathcal{I}} N_i^z(x_i)$, and (4) each $\succsim_i(z)$ is monotone (if $x_i, y_i \in \mathbb{R}_i^{nmp}(z)$ and $x_i > y_i$ then $x_i \succ_i(z)\ y_i$) and continuous (for each $x_i \in \mathbb{R}_i^{nmp}(z)$, the set $\{y_i \in \mathbb{R}_i^{nmp}(z) : y_i \succ_i(z)\ x_i\}$ is open relative to $\mathbb{R}_i^{nmp}(z)$).*

Instead of imposing Definition 9, it would be sufficient for our purposes to assume that each $\succsim_i(z)$ is strictly convex, monotone, continuous, differentiable, complete, and transitive, and that the $\succsim_i (z)$ taken together satisfy a 'no isolated communities' condition (see Arrow and Hahn (1971) or Mas-Colell (1985) for a detailed treatment).[8] To satisfy Definition 9, it will normally be necessary that each $\succsim_i$ does not display multiple supporting prices for at least one good.

Given an allocation z, let $mp(k, i) = 1$ if $\succsim_i$ displays multiple supporting prices for good k at z and $mp(k, i) = 0$ otherwise, and let $MP_k(z) = \max[\sum_{i \in \mathcal{I}} mp(k, i) - 1, 0]$.

Proposition 4 *If each i's behavioral welfare relation satisfies the convexity condition and x is a robust Pareto optimum that has a well-behaved conditional economy, then x is contained in a set of optima of dimension $I - 1 + \sum_{k=1}^{L} MP_k(x)$.*

[8] Explicitly, 'no isolated communities' requires that for every partition of $\mathcal{I}$ into $\mathcal{I}_1$ and $\mathcal{I}_2$ there exists a good k such that, for some $i \in \mathcal{I}_1$ and some $j \in \mathcal{I}_2$, $\succsim_i$ and $\succsim_j$ do not display multiple supporting prices for k at z. This condition, adapted from Smale (1974), ensures that utility can be continuously transferred among agents using only goods for which agents do not display multiple supporting prices.

7 Conclusion

If we build a behavioral welfare relation for each agent i where i's ranking of options requires unanimity among the frames at which i chooses, the set of Pareto optimal allocations can be very large. In the polar case where agents' behavioral welfare relations are supported by a maximal set of prices, the dimension of the Pareto optima will equal the dimension of the entire commodity space.

The problems we have considered in this paper do not plague all versions of behavioral welfare economics. If we could make interpersonal comparisons of utility across the agents that appear at different frames, then we could aggregate the frame-based preferences into a complete preference relation for any individual i. Kahneman et al. (1997) is an attempt in this direction. The present problems are in fact a particular difficulty of the Pareto criterion. Pareto optimality is widely seen to be an impractical policy guide due to the fact that policymakers never know with certainty the preferences of the individuals they are dealing with; to avoid making interpersonal comparisons of utility, policymakers must look for policy changes that do not harm any of the preference relations that individuals might potentially have. Such a test is so demanding that virtually any policy can be labeled Pareto optimal. The current difficulties are similar. When we use the behavioral welfare relation $\succsim_i$ to measure i's welfare, we are in effect insisting that no agent at any frame is harmed by a policy change. We thus get a similar expansion of the set of agents and a similar expansion of the set of optima.

8 Appendix: Proofs

Proof of Proposition 1. Let x be a Pareto optimum and define $B \equiv \sum_{i\in\mathcal{I}}\{y_i \in \mathbb{R}^L_+ : y_i \succsim_i x_i\}$. Suppose $\sum_{i\in\mathcal{I}} x_i \in \text{int}B$. Letting 1 be the index of the nonsatiated agent, there exists $\widehat{x}_1 \succ_1 x_1$ and, by reflexivity (see section 3) and the convexity condition, $\alpha x_1 + (1-\alpha)\widehat{x}_1 \succ_1 x_1$ for any $\alpha \in (0,1)$. Set α near enough to 1 to ensure that $z \equiv (2-\alpha)x_1 - (1-\alpha)\widehat{x}_1 + \sum_{i\neq 1} x_i \in B$. So there exists $(z_1, ..., z_I)$ such that $\sum_{i\in\mathcal{I}} z_i = z$ and $z_i \succsim_i x_i$ for all $i \in \mathcal{I}$. By the convexity condition, $\frac{1}{2}z_1 + \frac{1}{2}(\alpha x_1 + (1-\alpha)\widehat{x}_1) \succ_1 x_1$ and $\frac{1}{2}z_i + \frac{1}{2}x_i \succsim_i x_i$ for $i = 2, ..., I$. Given that $\frac{1}{2}[\alpha x_1 + (1-\alpha)\widehat{x}_1 + \sum_{i\neq 1} x_i] + \frac{1}{2}z = \sum_{i\in\mathcal{I}} x_i = e$,

x could not be a Pareto optimum.

Hence $\sum_{i\in\mathcal{I}} x_i \notin \text{int}B$. The convexity of B therefore implies there is a $p \neq 0$ such that $p \cdot y \geq p \cdot \sum_{i\in\mathcal{I}} x_i$ for any $y \in B$. For any $j \in \mathcal{I}$, let x'_j satisfy $x'_j \succsim_j x_j$. Then, since each $\succsim_i$ is reflexive, $x'_j + \sum_{i\neq j} x_i \in B$, and so $p \cdot x'_j \geq p \cdot x_j$. ■

Proof of Proposition 2. Let x be a robust Pareto optimum and let $p \in \bigcap_{i\in\mathcal{I}} \text{int}N_i(x_i)$. By the continuity of supporting prices, for any $x' \in Y$ sufficiently near x, $p \in \bigcap_{i\in\mathcal{I}} \text{int}N_i(x'_i)$. To show that x' is a Pareto optimum, suppose to the contrary that there is a $y \in Y$ with $y_i \succsim_i x'_i$ for all $i \in \mathcal{I}$ and $y_j \succ_j x'_j$ for some $j \in \mathcal{I}$. Then $p \cdot (y_i - x'_i) \geq 0$ for all $i \in I\!\!/$. If $p \cdot (y_j - x'_j) > 0$, then $\sum_{i\in\mathcal{I}} p \cdot (y_i - x'_i) > 0$, which is inconsistent with $\sum_{i\in\mathcal{I}} x'_i = \sum_{i\in\mathcal{I}} y_i = e$. To see that $p \cdot (y_j - x'_j) = 0$ cannot occur, suppose it does and label coordinates so that coordinate 1 has $y_j(1) \neq x'_j(1)$. Then, for all ε with $\|\varepsilon\|$ sufficiently small, $\widehat{p} = (p(1) + \varepsilon, (\lambda p(k))_{k=2,\ldots,L})$, where λ is set so that $\widehat{p} \in S$, will satisfy $\widehat{p} \in N_j(x'_j)$. By letting ε be < 0 if $y_j(1) > x'_j(l)$ and > 0 if $y_j(1) < x'_j(1)$, one may readily confirm that $\widehat{p} \cdot (y_j - x'_j) < 0$, a contradiction. ■

Proof of Proposition 3. Given Proposition 2, regarding the robust Pareto optima it remains only to show that (I) for an open dense set of economies, PO_r is nonempty. For the non-robust case, it remains to show that for an open dense set of economies (II) any non-robust optimum is the limit of a sequence of robust optima, thus implying $PO_{nr} \subset \partial PO_r$ (where PO_{nr} denotes the set of non-robust optima and ∂ denotes 'the boundary of'), and (III) the non-robust optima have measure 0.

To show that I, II, and III hold for an open and dense set of economies, we (1) define a finite-dimensional set of parameters δ, which will establish property I, (2) define a product space of the δ's and the endogenous variables (x, p) and a map for each agent i such that, for any δ, the map is a submersion onto M_i (3) use this map and the transversality theorem to show that generically the M_i intersect transversally, (4) show that transversal intersection of the M_i implies property II, (5) add an additional transversality argument to show that, for almost every allocation x, the $N_i(x_i)$ intersect transversally, to establish property III.

We first extend each M_i to remove the boundary on the x coordinate. Let $\mathbb{R}^{ex}$ be an open subset of the $L(I-1)$-dimensional affine subspace in $\mathbb{R}^{LI}$ that contains Y such that

$Y \subset \mathbb{R}^{ex}$. We extend each M_i to a C^1 manifold with boundary $M_i^{ex} \subset \mathbb{R}^{ex} \times S$ such that the projection of M_i^{ex} onto its x coordinate equals $\mathbb{R}^{ex}$.

(1) Given M_i^{ex} and $\delta = (\delta_1, ..., \delta_I) \in (0,1) \times \cdots \times (0,1) \equiv \Delta$, define

$$M_i(\delta) = \{(x, p+q) \in \mathbb{R}^{LI} \times S : q \in \beta(\delta_i, x) \text{ and } (x,p) \in M_i^{ex}\},$$

where $\beta(\delta_i, x)$ is the ball in $\mathbb{R}^L$ with center 0 and radius $r(\delta_i, x) = \frac{\delta_i(2- \text{ diam } N_i(x_i))}{2}$.[9] Also, let $N_i(x, \delta)$ denote $\{p \in S : (x,p) \in M_i(\delta)\}$; given $x_i \in \mathbb{R}^L_+$, define the convex sets $B_i(x_i) \equiv \{y_i \in \mathbb{R}^L_+ : y_i \succsim_i x_i\}$ and $\overline{B}_i(x_i) \equiv \{y_i \in B_i(x_i) : y_i$ is not an extreme point of $B_i(x_i)\}$; and finally, given some $p \in S$, define $H_p \equiv \{y_i \in \mathbb{R}^L_+ : p \cdot (y_i - x_i) \geq 0\}$. Define the reflexive relation $\succsim_{i,\delta}$ by setting

$$B_{i,\delta}(x_i) \equiv \{y_i \in \mathbb{R}^L_+ : y_i \succsim_{i,\delta} x_i\} = \{x_i\} \bigcup \left(\overline{B}_i(x_i) \bigcap \left(\bigcap_{p \in N_i(x,\delta)} H_p \right) \right)$$

for each $x_i \in \mathbb{R}^L_+$. Since $N_i(x_i) \subset N_i(x, \delta)$, the set of supporting prices at x for the relation $\succsim_{i,\delta}$ is $N_i(x, \delta)$. We have let the expansion of the N_i shrink to 0 as diam $N_i(x_i)$ approaches 2 so that diam $N_i(x, \delta) < 2$, thus ensuring that $N_i(x, \delta)$ remains the intersection of S and a convex cone, and hence that $B_{i,\delta}(x_i)$ is nonempty and therefore consistent with nonsatiation. That $\succsim_{i,\delta}$ is smooth follows from the fact that $\succsim_i$ is smooth and the fact that $r(\delta_i, x)$ is a C^1 function of x. Finally, our convexity condition must be satisfied by $\succsim_{i,\delta}$ since $B_{i,\delta}(x_i)$ is convex and since $y_i \sim_{i,\delta} x_i$ implies $y_i = x_i$.

If x is a Pareto optimum for the original economy with $\succsim_i$, $i \in \mathcal{I}$, Proposition 1 implies there is a p such that $p \in N_i(x_i)$ for all $i \in \mathcal{I}$. Hence for any $\delta \in \Delta$, x is a robust Pareto optimum for $\succsim_{i,\delta}$, $i \in \mathcal{I}$. And since we have assumed that the original economy has a Pareto optimum, for any $\delta \in \Delta$ the economy that results has a robust Pareto optimum. Since $\delta^n \to \delta$ implies that each $\succsim_{i,\delta^n}$ converges to $\succsim_{i,\delta}$, the economies for which PO_r is nonempty form a dense (and, self-evidently, open) set: property I is satisfied.

(2) Let Y^{ex} denote the $L(I-1)$-manifold formed by the intersection of $\mathbb{R}^{ex}$ and the affine subspace spanned by Y. Since, for any x and δ, $\partial N_i(x, \delta)$ is compact and boundaryless the ε-neighborhood theorem (see, e.g., Guillemin and Pollack (1974), 2.3) implies there exists a neighborhood of $\partial N_i(x, \delta)$ in, and open relative to, S, and a C^1

[9] For $A \subset \mathbb{R}^m$, diam $A = \sup_{x \in A, y \in A} \|x - y\|$.

submersion from this neighborhood to $\partial N_i(x,\delta)$ that is the identity on $\partial N_i(x,\delta)$; this function can be chosen to take each p to the $\widehat{p} \in \partial N_i(x,\delta)$ that minimizes $\|p-\widehat{p}\|$. By adjusting the proof of the ε-neighborhood theorem slightly, one may show that there is a set $\widehat{YS} \times \widehat{\Delta}$ in, and open relative to, $(Y^{ex} \times S) \times \Delta$, and a C^1 map $G_i : \widehat{YS} \times \widehat{\Delta} \longrightarrow \widehat{YS}$ such that (i) $\text{proj}_{Y^{ex}}\widehat{YS} \supset Y$, (ii) $\widehat{\Delta} \subset \mathbb{R}^I_+$ is an open rectangle with $0 \in \text{cl} \quad \widehat{\Delta}$, (iii) if $(x,p) \in \partial M_i(\delta)$ and $\delta \in \widehat{\Delta}$, then $(x,p) \in \widehat{YS}$, (iv) G_i maps $(x,p,\delta) \in \widehat{YS} \times \widehat{\Delta}$ to $(x,\widehat{p})$, where $\widehat{p} \in \partial N_i(x,\delta)$ minimizes $\|p-\widehat{p}\|$, and (v) for $\delta \in \widehat{\Delta}$, $g_i^\delta : \widehat{YS} \longrightarrow \widehat{YS}$ defined by $g_i^\delta(x,p) = G_i(x,p,\delta)$ is a submersion onto $\partial M_i(\delta)$. (Here and subsequently $\partial M_i(\delta)$ and $\text{int}M_i(\delta)$ will refer to the boundary and interior of the manifold $M_i(\delta)$.) Property (iv) is not essential, but it simplifies the calculation of a derivative in (3).

(3) To show, for any C^1 submanifold P of $\widehat{YS}$, that $G_i \pitchfork P$, it is sufficient for dim(Image $DG_i(x,p,\delta)$) to equal $\dim \widehat{YS} = L(I-1)+L-1$ for any $(x,p,\delta) \in \widehat{YS} \times \widehat{\Delta}$. Since for $\delta \in \widehat{\Delta}$, g_i^δ is a submersion onto $\partial M_i(\delta)$, and $\partial M_i(\delta)$ has dimension equal to $\dim(\widehat{YS})-1$, we have dim(Image $DG_i(x,p,\delta)) \geq \dim(\widehat{YS})-1$. Moreover, dim(Image $DG_i(x,p,\delta)) = \dim(\widehat{YS})$ if, for any $\delta \in \widehat{\Delta}$ and $(x,p) \in \partial M_i(\delta)$, Image $(DG_i(x,p,\delta))$ contains some direction not in $T_{x,p}\partial M_i(\delta)$. ($T_yA$ will denote the tangent bundle of a manifold A at y.) For $m \in T_{\widehat{YS}}$ given by $(0,p' \neq 0)$, where $p' \cdot p = 0$ for all $p \in \partial N_i(x,\delta)$, we have $m \perp T_{x,p}\partial M_i(\delta)$. Since $D_{\delta_i}G_i(x,p,\delta) = \frac{r(\delta_i,x)}{\|m\|}m$, $D_{\delta_i}G_i(x,p,\delta)m \neq 0$ and so m may serve as the additional direction.

For the submanifold $\partial M_j(\delta)$ of $\widehat{YS}$, $j \neq i$, the transversality theorem implies that the $\delta \in \widehat{\Delta}$ such that $g_i^\delta(x,p) \pitchfork \partial M_j(\delta)$ form a set $\Delta_{ij} \subset \widehat{\Delta}$ whose complement in $\widehat{\Delta}$ has 0 measure. Since g_i^δ is a submersion onto $\partial M_i(\delta)$, Image $Dg_i^\delta(x,p)$ coincides with Image $D\iota(M_i(\delta))$, where $\iota(M_i(\delta)) : \partial M_i(\delta) \hookrightarrow \widehat{YS}$ is the inclusion map of $\partial M_i(\delta)$ and so $\partial M_i(\delta) \pitchfork \partial M_j(\delta)$ for $\delta \in \Delta_{ij}$. Moreover, since Y is compact, the $\overline{\Delta}_{ij} \subset \widehat{\Delta}$ such that $\partial M_i(\delta) \pitchfork \partial M_j(\delta)$ on Y in addition to (i) having a 0-measure complement and (ii) containing 0 in its closure is also (iii) open. Call any subset of $\widehat{\Delta}$ with these three properties *generic*. Apply the same logic to any pair (k,l) of agents and take the intersection of the resulting $\overline{\Delta}_{kl}$, thus arriving at a generic set Δ^2. Since for any $\delta \in \Delta^2$ and any $i,j,l \in \mathcal{I}$, $\partial M_j(\delta) \pitchfork \partial M_l(\delta)$, $\partial M_j(\delta) \cap \partial M_l(\delta)$ is a C^1 manifold and hence there is a generic set $\overline{\Delta}_{i,j,l}$ such that

$$\partial M_i(\delta) \pitchfork (\partial M_j(\delta) \cap \partial M_l(\delta)) \tag{T}$$

holds for any $\delta \in \overline{\Delta}_{i,j,l}$. Consequently there is also a generic set Δ^3 such that T holds for any triple in $\mathcal{I}$ and $\delta \in \Delta^3$. Proceeding by induction we conclude there is a generic set Δ^I such that for any $\delta \in \Delta^I$ and any agent i and $\mathcal{I}_i \subset \mathcal{I} \backslash \{i\}$, $\partial M_i(\delta) \pitchfork \bigcap_{j \in \mathcal{I}_i} \partial M_j(\delta)$.

(4) For any $\delta \in \Delta^I$ and $(x, p) \in \bigcap_{i \in \mathcal{I}_i} M_i(\delta)$ (that is, an optimal x supported by p), consider the $\widehat{\mathcal{I}} \subset \mathcal{I}$ defined by $i \in \widehat{\mathcal{I}}$ if and only if $p \in \partial N_i(\delta)$ (that is, the agents for whom p in on the boundary of the prices that support x_i). Relabel agents so that $\widehat{\mathcal{I}} = \{1, ..., t\}$. Since

$$T_{x,p}\partial M_1(\delta) + T_{x,p}(\bigcap_{j=2}^{t} \partial M_j(\delta)) = \mathbb{R}^{\dim \widehat{YS}},$$

there must be $(x^1, p^1) \in (\text{int}M_1(\delta)) \bigcap (\bigcap_{j=2}^{t} \partial M_j(\delta))$ arbitrarily near (x, p). Similarly since

$$T_{x,p}\partial M_2(\delta) + T_{x,p}(\bigcap_{j=3}^{t} \partial M_j(\delta)) = \mathbb{R}^{\dim \widehat{YS}},$$

there must exist $(x^2, p^2) \in (\text{int}M_2(\delta)) \bigcap (\bigcap_{j=3}^{t} \partial M_j(\delta))$ arbitrarily near (x^1, p^1) and hence still in $\text{int}M_i(\delta)$. Proceeding by induction we conclude there is a $(x^t, p^t) \in \bigcap_{j=1}^{t} \text{int}M_j(\delta)$ that may be chosen to be arbitrarily near (x, p). Since $(x^t, p^t) \in \bigcap_{j=1}^{t} \text{int}M_j(\delta)$ implies $p^t \in \bigcap_{j=1}^{t} \text{int}_S N_j(x^t, \delta)$ (where int_S is the interior relative to S), we conclude that for any $\delta \in \Delta^I$ and any interior optimum $x >> 0$ there is a sequence of robust optima that converges to x. For a boundary x in contrast, it may be that any $x^t \notin Y$. To cover the non-robust boundary optima, we can apply (with no alterations) the logic from (2) onwards to an arbitrary coordinate subspace. Specifically, for each good k let $\mathcal{I}(k)$ denote an arbitrary strict subset of $\mathcal{I}$ (indicating the agents who do not consume k), let $\overline{\mathbb{R}} = \{x \in \mathbb{R}_+^{LI} : i \in \mathcal{I}(k) \text{ implies } x_i(k) = 0\}$, and let $\overline{Y} = Y^{ex} \cap \overline{\mathbb{R}}$. Letting $\overline{Y}$ take the place of Y^{ex}, we conclude, for δ in a generic set, that for any non-robust boundary optimum in $\overline{Y} \cap Y$ we may find a $x^t \in \overline{Y} \cap Y$ arbitrarily near x such that $\bigcap_{j=1}^{t} \text{int}_S N_j(x^t, \delta) \neq \varnothing$. Thus the property $PO_{nr} \subset \partial PO_r$ holds for a dense set of economies. Since $PO_{nr} \subset \partial PO_r$ follows from our manifolds having transversal intersection, the openness of the property $PO_{nr} \subset \partial PO_r$ follows as usual from the compactness of Y.

(5) Next we show that PO_{nr} has measure 0 for any $\delta \in \Delta^I$. Fix some $\delta \in \Delta^I$. For any $x \in Y$ and $i \in \mathcal{I}$, let $\widetilde{Y}_x$ and $\widetilde{S}_x$, with $\widetilde{Y}_x \times \widetilde{S}_x \subset \widehat{YS}$ and open relative to Y^{ew}

and S respectively, be such that if $(y,p) \in \partial M_i(\delta)$ and $y \in \widetilde{Y}_x$ then $p \in \widetilde{S}_x$. Since $\{\widetilde{Y}_x\}_{x\in Y}$ covers the compact set Y, we can restrict ourselves to some finite selection from $\{\widetilde{Y}_x\}_{x\in Y}$ that covers Y. Since $\delta \in \Delta^I$, we know that $g_i^\delta \pitchfork (\bigcap_{j\in\mathcal{I}_i} \partial M_j(\delta))$ for any i and $\mathcal{I}_i \subset \mathcal{I}\backslash\{i\}$. Hence by the transversality theorem the function $h_i^{\delta,y} : \widetilde{S}_x \longrightarrow \widehat{YS}$ defined by $h_i^{\delta,y}(p) = g_i^\delta(y,p)$ satisfies $h_i^{\delta,y} \pitchfork (\bigcap_{j\in\mathcal{I}_i} \partial M_j(\delta))$ for a.e. $y \in \widetilde{Y}_x$. For any of these y and any p such that $(y,p) \in \partial M_i(\delta) \bigcap (\bigcap_{j\in\mathcal{I}_i} \partial M_j(\delta))$, Image $Dh_i^{\delta,y}(p)$ has dimension equal to $\dim S - 1$ and consists only of directions $(0,\widehat{p})$ where $\widehat{p} \in T_n \partial N_i(y,\delta)$. For $p' \in T_n S$ such that $(0,p') \perp$ Image $Dh_i^{\delta,y}(p)$, it must be (given $h_i^{\delta,y} \pitchfork (\bigcap_{j\in\mathcal{I}_i} \partial M_j(\delta))$) that $(0,p') \in T_{(y,p)}(\bigcap_{j\in\mathcal{I}_i} \partial M_j(\delta))$. Hence $\partial N_i(y,\delta) \pitchfork (\bigcap_{j\in\mathcal{I}_i} \partial N_j(y,\delta))$. Given i and $\widetilde{Y}_x$, we can specify such a set of y in $\widetilde{Y}_x$, each with null complement in $\widetilde{Y}_x$, for any of the finite number of $\mathcal{I}_i \subset \mathcal{I}\backslash\{i\}$. Letting the finite selection from $\{\widetilde{Y}_x\}_{x\in Y}$ vary and then letting i vary and taking the intersection of the resulting finite number of sets, we conclude that any y outside of a null set of allocations has $\partial N_i(y,\delta) \pitchfork (\bigcap_{j\in\mathcal{I}_i} \partial N_j(y,\delta))$ for all i and $\mathcal{I}_i$. Hence any such y that is an optimum is a robust optimum. As in the previous paragraph, therefore, for any of these y that are optimal, $\bigcap_{i\in\mathcal{I}} \text{int}_S N_i(y,\delta) \neq \varnothing$. Openness of the property of PO_{nr} having measure 0 follows again from the compactness of Y. ■

Proof of Proposition 4. Let $\overline{x}$ be the robust optimum given in the Proposition. Condition 2 of robustness implies that for z in a neighborhood of $\overline{x}$ the coordinates k of z_i such that $\succsim_i$ displays multiple supporting prices for good k at z do not change as a function of z. Therefore, since our analysis is local, we can just refer to the coordinates for which $\succsim_i$ displays multiple supporting prices: for any $i \in \mathcal{I}$ and any $z \in \mathbb{R}^{LI}$ or $z_i \in \mathbb{R}^L$, z_{2i} will denote the coordinates of z_i for which $\succsim_i$ displays multiple supporting prices and z_{1i} will denote the remaining coordinates of z_i. Since $\succsim_i(z) = \succsim_i(y)$ whenever $z_{2i} = y_{2i}$, we write $\succsim_i(z_{2i})$ instead of $\succsim_i(z)$. Given $z \in \mathbb{R}^{LI}$ and $a = 1, 2$, let z_a denote $(z_{a1}, ..., z_{aI})$. Since $\overline{x}$ is a robust optimum, there is a $\overline{p}$ such that if $\succsim_i$ displays multiple supporting prices for k at $\overline{x}$ then $\overline{p} \in \text{int}_{S_k(\overline{p})}(N_i(\overline{x}_i) \cap S_k(\overline{p}))$. Since we have assumed that, for any z_2, $PO^{nmp}(z_2) \equiv \{z_1 : z_1$ is a Pareto optimum for $(e, \succsim_i(z_{2i})_{i\in\mathcal{I}})\}$ has dimension $I-1$, the set $\{x' \in Y : x'_1 \in PO^{nmp}(x'_2)\}$ has dimension $I - 1 + \sum_{k=1}^L MP_k(x)$. Hence it is sufficient to show, for any $x^n = (x_1^n, x_2^n) \longrightarrow \overline{x}$ such that $\sum_{i\in\mathcal{I}} x_i^n = \sum_{i\in\mathcal{I}} \overline{x}_i$ and $x_1^n \in PO^{nmp}(x_2^n)$ for each n that x^n is a Pareto optimum for all n sufficiently large. Let p_1 denote a generic element of $S^{nmp}(x)$, and given $p_1 \in S^{nmp}(x)$, let p_{1i} denote those coordinates k of p_1 such that $\succsim_i$ does not display multiple supporting prices for k. Given that $x_1^n \in PO^{nmp}(x_2^n)$, there is a sequence p_1^n where each $p_1^n \in \bigcap_{i\in\mathcal{I}} N_i^{x^n}(x_i^n)$. Define $P_i^n = \{z_i^n \in \mathbb{R}^L : p_{1i}^n \cdot z_{1i}^n = 0, z_{2i}^n = 0\}$. Since $x_1^n \in PO^{nmp}(x_2^n)$, $z_i^n \in P_i^n$ implies $x_i^n \succsim_i z_i^n + x_i^n$. Hence our convexity condition implies that $z_i^n + x_i^n$ cannot be interior to $\{y_i^n \in \mathbb{R}_+^L : y_i^n \succsim_i x_i^n\}$. Thus P_i^n cannot intersect the interior of $B_i(x_i^n) \equiv \{z_i^n \in \mathbb{R}_+^L : z_i^n + x_i^n \succsim_i x_i^n\}$. So there is a $q_i^n \in S$ such that $q_i^n \cdot z_i^n \geq 0$ if $z_i^n \in B_i(x_i^n)$ and $q_i^n \cdot z_i^n \leq 0$ if $z_i^n \in P_i^n$. In fact, $q_i^n \cdot z_i^n < 0$ cannot hold for any $\widetilde{z}_i^n \in P_i^n$ since then $q_i^n \cdot (-\widetilde{z}_i^n) > 0$ for $-\widetilde{z}_i^n \in P_i^n$. Thus $p_{1i}^n \cdot z_{1i}^n = 0$ implies $q_{1i}^n \cdot z_{1i}^n = 0$. By monotonicity, $q_{1i}^n >> 0$ and hence $\frac{1}{\|q_{1i}^n\|} q_{1i}^n = \frac{1}{\|p_{1i}^n\|} p_{1i}^n$. So, given that $\overline{p}_1 \in \bigcap_{i\in\mathcal{I}} N_i^{\overline{x}}(\overline{x}_i)$, the well-behaved conditional economies assumption implies $\frac{1}{\|q_{1i}^n\|} q_{1i}^n \longrightarrow \frac{1}{\|\overline{p}_{1i}\|}\overline{p}_{1i}$. Since $y_i^n \succsim_i x_i^n$ implies $q_i^n \cdot (y_i^n - x_i^n) \geq 0$, we have $q_i^n \in N_i(x_i^n)$. Given that $q_{1i}^n \longrightarrow \overline{p}_{1i}$ and $x_i^n \longrightarrow \overline{x}_i$, the robustness of $\overline{x}$ and condition 2 of robustness imply that $(p_{1i}^n, \overline{p}_{2i}) = (q_{1i}^n, \overline{p}_{2i}) \in N_i(x_i^n)$ (where $\overline{p}_{2i}$ are the coordinates k of $\overline{p}$ such that $\succsim_i$ displays multiple supporting prices for k). Since the goods for which any $\succsim_i$ does not display multiple supporting prices are a subset of the goods for which some $\succsim_j$ does not display multiple supporting prices, robustness and continuous supporting prices implies that $(p_1^n, \overline{p}_2) \in N_i(x_i^n)$ for all n sufficiently large and for all i. If y^n is a Pareto improvement over x^n, and j has both $y_j^n \succ_j x_j^n$ and $y_{2j}^n \neq x_{2j}^n$, then we can follow the proof of Proposition 2 to show that $(p_1^n, \overline{p}_2) \cdot (y_j^n - x_j^n) > 0$ which, combined with $(p_1^n, \overline{p}_2) \cdot (y_i^n - x_i^n) \geq 0$ for all $i \in \mathcal{I}$, implies that $y^n \notin Y$. If on the other hand j has both $y_j^n \succ_j x_j^n$ and $y_{1j}^n \neq x_{1j}^n$ then, as in the standard proof of the first welfare theorem, continuity rules out the possibility $(p_1^n, \overline{p}_2) \cdot (y_j^n - x_j^n) = 0$; hence $(p_1^n, \overline{p}_2) \cdot (y_j^n - x_j^n) > 0$, again violating feasibility. ■

References

[1] Arrow, K., and F. Hahn, 1971, *General competitive analysis*, North-Holland, Amsterdam.

[2] Bernheim, B. D., and A. Rangel, 2007, 'Toward choice-theoretic foundations for behavioral welfare economics,' *American Economic Review* 97: 464-470.

[3] Bernheim, B. D., and A. Rangel, 2008, 'Beyond revealed preference: choice theoretic foundations of behavioral welfare economics,' forthcoming, *Quarterly Journal of Economics.*

[4] Bewley, T., 1986, 'Knightian decision theory: part 1,' *Cowles Foundation Discussion Paper* 807, New Haven CT.

[5] Billot, A., A. Chateauneuf, I. Gilboa, and J.-M. Tallon, 2000, 'Sharing beliefs: between agreeing and disagreeing,' *Econometrica* 68: 685–694.

[6] Bonnisseau, J.-M., and B. Cornet, 1988, 'Valuation equilibrium and Pareto optimum in non-convex economies,' *Journal of Mathematical Economics* 17: 293-308.

[7] Dana, R. A., 2004, 'Market behavior when preferences are generated by second-order stochastic dominance,' *Journal of Mathematical Economics* 40: 619–639.

[8] Fon, V., and Y. Otani, 1979, 'Classical welfare theorems with non-transitive and non-complete preferences,' *Journal of Economic Theory* 20: 409-418.

[9] Guillemin, V., and A. Pollack, 1974, *Differential topology*, Prentice-Hall, Englewood Cliffs.

[10] Kahneman, D., J. Knetsch, and R. Thaler, 1991, 'The endowment effect, loss aversion, and status quo bias: anomalies,' *Journal of Economic Perspectives* 5: 193-206.

[11] Kahneman, D., P. Wakker, and R. Sarin, 1997, 'Back to Bentham? Explorations of experienced utility,' *Quarterly Journal of Economics* 112: 375-406.

[12] Laibson, D., 1997, 'Golden eggs and hyperbolic discounting,' *Quarterly Journal of Economics* 112: 443-477.

[13] Mandler, M., 2004a, 'Status quo maintenance reconsidered: changing or incomplete preferences?' *Economic Journal* 114: 518-535.

[14] Mandler, M., 2004b, 'Welfare economics with status quo bias: a policy paralysis problem and cure,' mimeo, Royal Holloway College, University of London.

[15] Mandler, M., 2005, 'Incomplete preferences and rational transitivity of choice,' *Games and Economic Behavior* 50: 255–277.

[16] Mas-Colell, A., 1985, *The theory of general economic equilibrium: a differentiable approach*, Cambridge University Press: Cambridge.

[17] O'Donoghue, T., and M. Rabin, 1999, 'Doing it now or later,' *American Economic Review* 89: 103-24

[18] Rigotti, L. and C. Shannon, 2005, 'Uncertainty and risk in financial markets,' *Econometrica* 73: 203-243.

[19] Salant, Y. and A. Rubinstein, 2008, '(A, f) choice with frames,' forthcoming, *Review of Economic Studies*.

[20] Smale, S., 1974, 'Global analysis and economics V,' *Journal of Mathematical Economics* 1: 213-221.

[21] Tversky, A., and D. Kahneman, 1991, 'Loss aversion in riskless choice: a reference-dependent model,' *Quarterly Journal of Economics* 106: 1039-61.