

Beyond Revealed Preference: Toward Choice-Theoretic Foundations for Behavioral Welfare Economics

B. Douglas Bernheim and Antonio Rangel

December, 2006

Interest in behavioral economics has grown in recent years, stimulated largely by accumulating evidence that the standard model of consumer decision-making provides an inadequate positive description of human behavior. Behavioral models are increasingly finding their way into policy evaluation, which inevitably involves welfare analysis. The welfare issues raised by behavioral economics have not received nearly as much attention as the positive issues. No consensus concerning the appropriate standards and criteria for behavioral welfare analysis has yet emerged.

The most common mode of behavioral welfare analysis involves distinguishing between “decision utility,” which provides an as-if representation of choices, possibly by invoking unconventional assumptions concerning preferences, and “true” or “experienced” utility, which is viewed as the proper measure of well-being. This approach forces one to take a stand on the nature of true utility. But the objective basis for making any assumptions about true utility is, at best, obscure. Evidence of incoherent choice patterns, coupled with the absence of a scientific foundation for assessing true utility, has led some to conclude that behavioral economics must embrace fundamentally different normative principles than standard economics (see, e.g., Sugden [***]).

This paper summarizes our effort to develop a unified framework for behavioral welfare economics (see Bernheim and Rangel [2007]) – one that can be viewed as a natural extension

of standard welfare economics. We reject the notion that it is necessary to recover “true utility;” indeed, our approach does not even require one to regard this concept as meaningful. Standard welfare analysis is about choice, not utility or preferences. In its simplest form, it instructs the planner to respect the choices an individual would make for himself. The guiding normative principle is simply an extension of the libertarian deference to freedom of choice – standard welfare analysis adopts the view that it is better to give a person the thing he would choose for himself rather than something that someone else would choose for him. Significantly, the application of this principle does not require all choices to be consistent. Though the guidance provided by choice data may be ambiguous in some circumstances, it may nevertheless be unambiguous in others. As we show, this partially ambiguous guidance provides a choice-theoretic foundation for rigorous welfare analysis. Moreover, one can carry out this analysis using straightforward generalizations of standard tools, such as compensating and equivalent variation, consumer surplus, and Pareto optima.

We also propose an agenda for refining our welfare criterion based on non-choice data. Significantly, we do not propose *substituting* non-choice data, or any external judgment, for choice data. Rather, we adhere to the principle that the planner’s objective should be to select an alternative that the individual would select for himself *under some circumstances*. Non-choice data are potentially valuable because they may provide important information concerning *which* choice circumstances are most relevant for policy.

1 Preliminaries

1.1 Choice situations and choices

We will use \mathbb{X} to denote the set of all possible choice objects. Elements of \mathbb{X} could be standard objects such as state-contingent consumption bundles, or lotteries over such bundles, but these objects could also have non-standard features (e.g., as in Caplin and Leahy [2001]). .

A *standard choice situation* (SCS) consists of a constraint set $X \subseteq \mathbb{X}$, representing the set of objects from which the individual can choose according to the objective information

at his disposal. We are concerned with making welfare judgments based on behavior for some particular domain of standard choice situations, \mathcal{X} .

To accommodate behavioral anomalies, we introduce the notion of a *generalized choice situation* (GCS), G . This consists of a standard choice situation, X , and a set of ancillary conditions, d . Thus, $G = (X, d)$. We will use \mathcal{G} to denote the set of generalized choice situations of potential interest. Behavioral is associated with a choice correspondence $C : \mathcal{G} \rightarrow \mathbb{X}$; for any $G \in \mathcal{G}$, any $x \in C(G)$ is an action that the individual is willing to choose.

Research in behavioral economics has drawn attention to specific types of ancillary conditions, such as the point in time at which the decision is made, the manner in which information is presented, and the nature of the status quo or default option. As a general matter, it is difficult to draw a bright line between the characteristics of the objects in X and the ancillary conditions d (which, in principle, one could also view as a characteristic of the objects in the choice set). Our analytic framework is applicable regardless of how one draws this distinction (see Bernheim and Rangel [2007] for a discussion of some possible criteria).

Standard economics proceeds from the assumption that choice is invariant with respect to ancillary conditions. Positive behavioral economics challenges this basis premise. Documentation of a behavioral anomaly usually involves identifying some SCS, X , along with two ancillary conditions, d' and d'' , for which $C(X, d') \neq C(X, d'')$. This is sometimes called a *preference reversal*, but in the interests of greater precision we will call it a *choice reversal*.

1.2 Positive versus normative analysis

Usually, choice data are not available for all elements of \mathcal{G} , but rather for elements of some set $\mathcal{H} \subset \mathcal{G}$. The objective of positive economic analysis is to extend the choice correspondence C from observations on \mathcal{H} to the entire set \mathcal{G} . In standard economics, this is accomplished by defining preferences over \mathbb{X} , estimating these preferences with choice data for the opportunity sets in \mathcal{H} , and using these estimated preferences to infer choices for opportunity sets in $\mathcal{G} \setminus \mathcal{H}$.

The objective of normative economic analysis is to evaluate outcomes. Typically, we

evaluate outcomes at the level of the individual, and then aggregate. For choice-based normative criteria, we allow the individual’s choices to govern these evaluations. The fundamental problem of behavioral welfare economics is to identify appropriate criteria for evaluating alternatives when the individual’s choices fail to provide clear guidance.

In conducting choice-based normative analysis, we take as given the individual’s choice correspondence, C , defined on \mathcal{G} rather than \mathcal{H} . Preferences and utility functions, which are constructs used both in standard theory and in behavioral economics to extend C from \mathcal{H} to \mathcal{G} , are therefore positive tools, not normative tools. In a behavioral setting, these constructs cannot meaningfully reconcile inconsistencies; they can only reiterate the information contained in the extended choice correspondence C (both the observed choices and the inferred choices). Thus, one cannot resolve normative puzzles by identifying classes of preferences that rationalize apparently inconsistent choices, as in Gul and Pesendorfer [2001] (see Bernheim and Rangel [2007] for an analysis of normative ambiguities with “temptation preferences”).

2 A Framework for Behavioral Welfare Analysis

In the standard approach to normative analysis, one evaluates individual welfare by applying a binary relation, R , defined over the elements of \mathbb{X} . While we refer to this relation as “revealed preference,” it is simply a summary of what is chosen from various SCSs. Specifically, xRy means that if x and y belong to the set X , and if $y \in C(X)$, then $x \in C(X)$. Under standard assumptions, R is an ordering. When we use this ordering to conduct welfare analysis, we are simply asking what an individual or individuals would choose. For example, the compensating variation associated with some change in the economic environment equals the smallest payment that would induce the individual to choose the change. A social alternative $x \in X$ is a Pareto optimum in X if there is no other alternative in X that all individuals would choose over x .

Findings in behavioral economics suggest that we cannot summarize all choices with a

consistent preference ordering. Instead, choice evidence is sometime ambiguous. Fortunately, we can live with this ambiguity. It is still possible to construct a binary relation, R , based on *unambiguous* comparisons. A direct generalization of the standard approach would simply take xRy to mean that if x and y belong to the set X , and if $y \in C(X, d)$, then $x \in C(X, d)$. Because of choice reversals and other ambiguities, this relation will in general be neither complete nor transitive. However, as we show below, this approach provides a basis for conducting coherent welfare analysis, and allows us to generalize concepts such as compensating variation and Pareto optima.

2.1 Individual welfare orderings

For any $x, y \in \mathbb{X}$, we will say that xRy if, whenever x and y are available, x is sometimes chosen, and y is never chosen unless x is as well. Formally, (i) there exists some $(X, d) \in \mathcal{G}$ with $\{x, y\} \subseteq X$ such that $x \in C(X, d)$, and (ii) there does not exist any $(X, d) \in \mathcal{G}$ with $\{x, y\} \subseteq X$ such that $y \in C(X, d)$ and $x \notin C(X, d)$. (If $\{x, y\} \in \mathcal{X}$, then the second requirement implies the first.) Note that xRx .

As usual, we can define xPy as xRy and $\sim yRx$. This means that, whenever x and y are available, sometimes x is chosen but not y , and otherwise either both or neither are chosen. Likewise, we can define xIy as xRy and yRx . This means that, whenever x is chosen, so is y , and vice versa.

Finally, we will say that xP^*y if, whenever x and y are available, sometimes x is chosen but not y , and otherwise neither are chosen. Formally, (i) for all $(X, d) \in \mathcal{G}$ with $\{x, y\} \subseteq X$, we have $y \notin C(X, d)$, and (ii) for some $(X, d) \in \mathcal{G}$ with $\{x, y\} \subseteq X$, we have $x \in C(X, d)$.

In general, R , P , and P^* need not be orderings. For example, if $C(\{x, y\}, d') = \{x\}$ and $C(\{x, y\}, d'') = \{y\}$, then we have *neither* xRy *nor* yRx , so R is not complete. Moreover, R , P , and P^* needn't be transitive. For example, if choice does not depend on ancillary conditions, and if we have $C(\{x_1, x_2\}) = x_1$, $C(\{x_2, x_3\}) = x_2$, $C(\{x_3, x_1\}) = x_3$, and $C(\{x_1, x_2, x_3\}) = \{x_1, x_2, x_3\}$, then $x_1Px_2Px_3Px_1$.

As we will see, these binary relations nevertheless provide a basis for meaningful welfare

analysis.

2.2 Individual welfare optima

There are two natural criteria for determining whether a choice is improvable. We will say that it is possible to *strictly improve* upon a choice $x \in X$ if there exists $y \in X$ such that yP^*x – in other words, there is an alternative that is unambiguously chosen over x . When a strict improvement is impossible, we say that x is a *weak individual welfare optimum*. It is possible to *weakly improve* upon a choice $x \in X$ if there exists $y \in X$ such that yPx – in other words, there is an alternative that is sometimes chosen over x , and that x is never chosen over (except in the sense that both could be chosen). When a weak improvement is impossible, we say that x is a *strict individual welfare optimum*.

When is $x \in X$ an individual welfare optimum? The following simple results (which we state without proof) assume that \mathcal{G} includes all pairwise comparisons (that is, for all $a, b \in \mathbb{X}$, there is some d_{ab} such that $(\{a, b\}, d_{ab}) \in \mathcal{G}$).

Fact 1: x is a weak individual welfare optimum in X if and only if for each $y \in X$ (other than x), there is some GCS for which x is chosen with y present (y may be chosen as well).

This result has an immediate corollary: if x is chosen for some GSC involving X , then x is a weak individual welfare optimum in X . Notice that this corollary guarantees the existence of weak individual welfare optima.

Fact 2: x is a strict individual welfare optimum in X iff for each $y \in X$ (other than x), either x is chosen and y is not for some GCS with y present, or there's no GCS for which y is chosen and x is not with x present.

This result also has an immediate corollary: if x is the unique choice for some GSC involving X , then x is a strict individual welfare optimum in X . A strict individual welfare optimum may not exist (see the example given at the end of the last section).

2.3 Relationship to multi-self Pareto optima

Our notion of an individual welfare optimum is related to the idea of a multi-self Pareto optimum. Suppose in particular that the set of GCSs is the Cartesian product of the set of SCSs and a set of ancillary conditions (that is, $\mathcal{G} = \mathcal{X} \times D$, where $d \in D$). Also imagine that, for each $d \in D$, choices follow standard axioms, and can be represented by a preference ranking R_d . If one imagines that each ancillary condition activates a different “self,” then one can conduct welfare analysis by examining multi-self Pareto optima. Under the stated conditions, a weak multi-self Pareto optimum corresponds to a weak individual welfare optimum (as we have defined it), and a strict multi-self Pareto optimum corresponds to a strict individual welfare optimum. For these narrow settings, one can therefore view our approach as a way to justify the multi-self Pareto criterion without relying on untested and questionable psychological assumptions. Our approach is also more general, in that it does not require the assumptions stated at the outset of this paragraph.

The multi-self Pareto criterion has been used primarily in the literature on quasi-hyperbolic discounting, where it is applied to an individual’s many time-dated “selves” (see, e.g., Laibson, Repetto, and Tobacman [1998] and Bhattacharya and Lakdawalla [2004]). Ironically, our framework does *not* justify the multi-self Pareto criterion for quasi-hyperbolic consumers, because the conditions listed in the previous paragraph are not satisfied; see Bernheim and Rangel [2007] for a full analysis of this environment.

2.4 Equivalent and compensating variation

The concepts of equivalent and compensating variation are central to applied welfare economics. Fortunately, they have natural counterparts within our framework. Here, we will focus on compensating variation; our treatment of equivalent variation is analogous. (See Bernheim and Rangel [2007] for some alternatives to the notions of compensated and equivalent variation defined here.)

We will write the individual’s SCS as $X(\alpha, m)$, where α is a vector of parameters, and m

is the level of compensation. Let α_0 be the initial parameter vector, d_0 the initial ancillary conditions, and $(X(\alpha_0, 0), d_0)$ the initial GCS. We will consider a change in parameters to α_1 , coupled with some level of compensation, as well as (potentially) a change in ancillary conditions that could in principle depend on the compensation level. We write the new GCS as $(X(\alpha_1, m), d(m))$. One natural possibility – but certainly not the only one – is to take $d(m) = d'$ for some fixed d' . This allows us, for example, to evaluate compensating variations for fixed changes in prices, ancillary conditions, or both.

We will define the compensating variation relative to particular selections from the choice correspondence. (Where the choice is not unique, one can imagine that it is tied down by an equilibrium condition, e.g. that supply must equal demand.) Accordingly, we assume the individual selects $x_0 \in C(X(\alpha_0, 0), d_0)$, and $x(m) \in C(X(\alpha_1, m), d(m))$.

In defining the compensating variation, we encounter an ambiguity concerning the standard of compensation: do we consider compensation sufficient when $x(m)$ is always weakly chosen over x_0 , or when x_0 is not always weakly chosen over $x(m)$? This ambiguity is an essential feature of welfare evaluations with inconsistent choice. Accordingly, we define two notions of compensating variation:

$$m^{CV-A} = \inf_m x(m)Rx_0$$

$$m^{CV-B} = \sup_m x_0Rx(m)$$

In other words, m^{CV-A} is the smallest level of m such that $x(m)$ is weakly chosen over x_0 . In contrast, m^{CV-B} is the largest level of m such that x_0 is weakly chosen over $x(m)$.

We will illustrate the application of these concepts by discussing the measurement of consumer surplus.

2.5 Consumer surplus

For simplicity, we will examine a case where the individual consumes two goods, x and y . Suppose that positive analysis delivers the following utility representation (which involves

no income effects, so that Marshallian consumer surplus would be valid in the standard framework):

$$U(x, y \mid d) = x + dv(y)$$

with v is strictly increasing, differentiable, and strictly concave. (See Bernheim and Rangel [2007] for a more general treatment of consumer surplus.) Thus, for any given d , the inverse demand curve for y is given by $p = dv'(y)$, where p is the price of y . Notice that the ancillary condition, $d \in [d_L, d_H]$, simply shifts the weight attached to $v(y)$. This might, for example, represent the type of “coherent arbitrariness” documented by Ariely, Loewenstein, and Prelec [2003].

Let M denote the consumer’s initial income. Consider a change in the price of y from p_0 to p_1 , along with a change in ancillary conditions from d_0 to d_1 (we mean to include cases where either $p_0 = p_1$ or $d_0 = d_1$). Let y_0 denote the amount of y purchased at price p_0 , and let y_1 denote the amount purchased at price p_1 ; assume that $y_0 > y_1$. Since there are no income effects, y_1 will not change as the individual is compensated (holding the ancillary condition fixed).

Now let’s calculate CV-A. We wish to find the smallest m such that $(y_1, M - p_1y_1 + m)R(y_0, M - p_0y_0)$. It is straightforward to show that the solution is

$$m^{CV-A} = [p_1 - p_0]y_1 + \int_{y_1}^{y_0} [d_H v'(y) - p_0]dy$$

Through similar reasoning, one can show that

$$m^{CV-B} = [p_1 - p_0]y_1 + \int_{y_1}^{y_0} [d_L v'(y) - p_0]dy$$

Figures 1(a) and 1(b) illustrate, respectively, m^{CV-A} and m^{CV-B} . Notice that these values bracket standard consumer surplus. Moreover, as the range of possible ancillary conditions narrows, they converge to standard consumer surplus. This underscores the fact that the standard framework is a special case of the framework considered here. Moreover, it also implies that, when inconsistencies are minor (that is, $d_H - d_L$ is small), the ambiguity in welfare, as measured by the difference between m^{CV-A} and m^{CV-B} , is small.

2.6 Generalized Pareto optima

Next we turn to environments with many individuals, and formulate a generalization of Pareto efficiency. Suppose there are N individuals indexed $i = 1, \dots, N$. Let \mathbb{X} denote the set of all conceivable social choice objects, and let X denote the set of feasible objects. Let C_i be the choice function for individual i , defined over \mathcal{G}_i (where the subscript reflects the possibility that the set of ancillary conditions may differ from individual to individual). These choice functions induce the relations R_i over \mathbb{X} .

We will say that x is a *weak generalized Pareto optimum* in X if there exists no $y \in X$ with yP_i^*x for all i . We will say that x is a *strict generalized Pareto optimum* in X if there exists no $y \in X$ with yR_ix for all i , and yP_i^*x for some i . (In Bernheim and Rangel [2007], we introduce some other variants of these notions.)

Since strict individual welfare optima do not always exist, we cannot guarantee the existence of strict generalized Pareto optima with a high degree of generality. However, we *can* guarantee the existence of a weak generalized Pareto optimum (for any compact set X) under a mild continuity condition (see Bernheim and Rangel [2007]).

2.7 The efficiency of competitive equilibria

To illustrate the usefulness of these concepts, we have provided a generalization of the first welfare theorem (see Bernheim and Rangel [2007]). Specifically, we consider a production economy consisting of N consumers, F firms, and K goods. The economy is standard in all respects, except that consumer i 's behavior is governed by a general choice correspondence mapping budget sets and ancillary conditions into sets of consumption vectors. We make one simple assumption (akin to non-satiation) with respect to consumer behavior: if $x^n > w^n$ (where $>$ indicates a strict inequality for every good), then consumer n never chooses w^n when x^n is available.

A *behavioral competitive equilibrium* involves a price vector, $\hat{\pi} = (\hat{\pi}^1, \dots, \hat{\pi}^K)$, along with a vector of ancillary conditions $\hat{d} = (\hat{d}^1, \dots, \hat{d}^N)$, that clear all markets. Though behavioral

competitive equilibria may not exist, those that do exist are necessarily efficient:

Theorem: The allocation in any behavioral competitive equilibrium is a strict generalized Pareto optimum.

Just as it is possible to demonstrate the efficiency of behavioral competitive equilibrium with perfectly competitive markets, it is also possible to demonstrate the inefficiency of equilibria in the presence of sufficiently severe market failures (see Bernheim and Rangel [2007] for an example). It is also worth emphasizing that even a perfectly competitive equilibrium may be inefficient when judged by a refined welfare relation, after officiating choice conflicts, as described in sections 3. This observation alerts us to the fact that, in behavioral economies, there is a new class of potential market failures.

3 Refinements

3.1 The logic of refinements

In any particular context, the relation R may not be very discerning, which means that the notion of an individual welfare optimum may not be very demanding. In this section, we consider the possibility that one might refine the relation R by altering the data used to construct it. Most obviously, one can add choice data (by creating new GCSs, expanding the domain \mathcal{G}), or delete data (by ignoring certain GCSs, reducing the domain \mathcal{G}). There is also the possibility of reinterpreting choice data, which we mention briefly below.

We will say that R' is *coarser* than R if $xR'y$ implies xRy . When R' is coarser than R , we say that R is *finer* than R' . Subject to a technical qualification, the addition of data (that is, the expansion of \mathcal{G}) makes R weakly coarser, while the elimination of data (that is, the reduction of \mathcal{G}) makes R weakly finer. Thus, to usefully refine the welfare relations, one must either eliminate or reinterpret data. Accordingly, if there is one GCS in which x is chosen over y , and another in which y is chosen over x , we look for objective criteria that might allow us to officiate between these GCSs, with the object of discarding

or reinterpreting one of them. We can then construct a new welfare relation, R' , based on the revised choice correspondence, which may be finer than R and contain fewer welfare optima.

How might we officiate between conflicting GCSs? One seemingly natural possibility, which we call “self-officiation,” is to officiate based on choices: if the individual makes conflicting choices for two GCSs, $G_1 = (X, d_1)$ and $G_2 = (X, d_2)$, simply allow him to choose between these GCSs. This, however, creates another GCS, call it $G_3 = (X, d_3)$. Since the expansion of \mathcal{G} to include G_3 does not usefully refine either the welfare relation or the sets of welfare optima, we have an “irresolvability principle”: new choices cannot resolve normative ambiguities associated with existing choice conflicts. To officiate, we must therefore rely on non-choice data..

3.2 Refinements based on information processing

When we say that an individual’s standard choice situation is X , we mean that, based on all of the objective information that is available to him, he is actually choosing among elements of X . In standard economics, we use this objective information to reconstruct X , and then infer that he prefers his chosen element to all the unchosen elements of X . But what if he fails to use all of the information available to him, or uses it incorrectly? What if the objective information available to him implies that he is actually choosing from the set X , while in fact he believes he is choosing from some other set, Y ? In that case, should a planner nevertheless mimic his choice when making a selection from X ? Not in our view.

Why would the individual believe himself to be choosing from some set, Y , when in fact, according to the available objective information, he is choosing from the set X ? His attention may focus on some small subset of X . His memory may fail to call up facts that relate choices to consequences. He may forecast the consequences of his choices incorrectly. He may have learned from his past experiences more slowly than the objective information would permit. Thus, the operations of these cognitive processes pertain to the question of whether, at the moment of choice, he appreciates that he is choosing from X .

In principle, if we understand the individual’s cognitive processes sufficiently well, we may be able identify his perceived choice set Y , and reinterpret the choice as pertaining to the set Y rather than to the set X . We refer to this process as “deconstructing choices.” While it may be possible to accomplish this in some instances (see, e.g., Koszegi and Rabin [2007]), we suspect that, in most cases, this is beyond the current capabilities of economists, cognitive neuroscientists, and psychologists. We nevertheless submit that there are circumstances in which non-choice evidence can reliably establish the existence of a significant *discrepancy* between the actual choice set, X , and the perceived choice set, Y . This occurs, for example, in circumstances where it is known that attention wanders, memory fails, forecasting is naive, and/or learning is inexplicably slow. In these instances, we say that the GCS is suspect.

We propose using non-choice evidence to officiate between conflicting choice data by deleting suspect GCSs. Thus, for example, if someone chooses x from X under condition d' where he is likely to be distracted, and chooses y from X under condition d'' where he is likely to be focused, we would delete the data associated with (X, d') before constructing the welfare relation R . In effect, we take the position that (X, d'') is a better guide for the planner than (X, d') . Even with the deletion of choice data, the welfare relation R may remain ambiguous in many cases due to other unresolved choice conflicts, but it nevertheless becomes (weakly) finer, and the sets of welfare optima grow (weakly) smaller.

What types of non-choice evidence might one use to determine the circumstances in which internal information processing systems work well, and the circumstances in which they work poorly? Evidence from neuroscience concerning the functioning of various cognitive processes can potentially shed light on the operation of processes governing attention, memory, forecasting, and learning. This evidence can provide an objective basis for determining whether a particular choice situation is suspect. For example, if memory is shown to function poorly under certain environmental conditions, GCSs that are associated with those conditions, and that require factual recall, are suspect. Our work on addiction (Bernheim and Rangel [2004]) provides a more elaborate illustration. Citing evidence from

neuroscience, we argue that the habitual use of addictive substances causes specific information processing systems to malfunction under identifiable circumstances. The choices made in these circumstances are therefore suspect, and welfare evaluations should be guided by choices made in other circumstances. More generally, these observations define a *normative* agenda for the emerging field of neuroeconomics.

In many situations, simpler forms of evidence may suffice. If, for example, an individual characterizes a choice as a mistake on the grounds that he neglected or misunderstood information, this may provide a compelling basis for declaring the choice suspect. Measures of affect could conceivably prove useful in detecting a mistake. Other considerations, such as the complexity of a GCS, could also come into play.

3.3 Other criteria for officiation

In Bernheim and Rangel (2007), we discuss other potential criteria for refining the welfare relation. One natural candidate is preponderance. In other words, if someone ordinarily chooses x over y (that is, in almost all choice situations where both are available and one is chosen), and rarely chooses y over x , it might be appropriate to disregard the exceptions and follow the rule. Another possibility is to rely on coherent patterns, rather than isolated anomalies, to guide welfare judgments. ***say something more? conclude briefly?***

***references