The Foerder Institute for Economic Research
at Tel Aviv University

The Eitan Berglas School of Economics

מכון למחקר כלכלי על שם
ד"ר ישעיהו פורדר
על יד אוניברסיטת תל אביב

בית-הספר לכלכלה ע"ש איתן ברגלס

עמותה רשומה

What Were You Thinking? Revealed Preference Theory as Coherence Test

Itzhak Gilboa and Larry Samuelson

Working Paper No. 4-2020

The Foerder Institute for Economic Research
and
The Sackler Institute of Economic Studies

מכון פורדר, אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409255 | פקס: 03-6405815
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409255 | Fax 972-3-6405815 | email: foerder@tauex.tau.ac.il

What Were You Thinking?
Revealed Preference Theory as Coherence Test*

Itzhak Gilboa† and Larry Samuelson‡

January 13, 2020

Abstract

Theory can be used to test the logic of intuitive decision making—one may ask whether a given set of decisions can be justified by a decision theoretic model in a given class. Indeed, in principal-agent settings, such justifications may be required—a manager of an investment fund may be asked what beliefs she had in mind when making financial decisions for her clients, or when evaluating assets and liabilities. While such a question is formally equivalent to a revealed preference question, our motivation suggests different assumptions about observable data. In this paper we assume that states and utilities are observable, and ask which collections of uncertain-act evaluations can be simultaneously justified by a single probability (for a Bayesian agent) and by a single set of probabilities (for a maxmin expected utility agent). We use a linear-programming-based argument to develop characterization results for each case.

Contents

*Gilboa gratefully acknowledges ISF Grant 1077/17, the Investissements d'Avenir ANR -11- IDEX-0003 / Labex ECODEC No. ANR - 11-LABX-0047, the AXA Chair for Decision Sciences at HEC and the Foerder Institute at Tel-Aviv University. We are grateful to Rossella Argenziano for comments.
†Tel-Aviv University and HEC, Paris. tzachigilboa@gmail.com
‡Department of Economics, Yale University. Larry.Samuelson@yale.edu

What Were You Thinking?
Revealed Preference Theory as Coherence Test

January 13, 2020

1 Introduction

1.1 Revealed Preference and Coherence

The initial forays into revealed preference theory (e.g., Samuelson (1938) and Houthakker (1950) were designed to clarify the conceptual foundations of models of decision making. These papers argued that an analyst could dispense with Edgeworth's (1881, Appendix III) hedonimeter and still reasonably model a person as maximizing utility, believing only that the person's choices satisfied certain consistency conditions. More generally, revealed preference theory can be interpreted as identifying the empirical content of models of decision making, facilitating the evaluation and application of such models (e.g., Andreoni and Miller (2002)). These are positive interpretations of revealed preference theory, geared toward "understanding what economic models say about the observable world" (Chambers and Echenique (2016,p. xiii)).

Revealed preference analyses can also be given a normative interpretation. An individual may wonder whether her decisions admit an expected utility representation because she aspires to conform to the underlying axioms (see, for example, Gilboa (2009, Section 6.3.3). An auditor may wonder whether there is a consistent view of the relevant uncertainty under which the various bets made by a fund manager can be justified. A venture capitalist may have similar concerns about the various aspects of an entrepreneur's business plan, or a voter about the cost and revenue projections in the various policy proposals of a political candidate. A regulator may wonder whether there is a common belief under which various parties can all benefit from a proposed transaction (Gilboa, Samuelson, and Schmeidler (2014)). We refer to these applications of revealed preference theory as *coherence tests*.

The "coherence test" question is formally identical to the revealed preference one—which data are compatible with a given class of models? However, there are two conceptual differences between the two interpretations. We have noted that the revealed preference interpretation is descriptive, while the coherence test is normative—in the former the modeler has no control over the data, whereas in the latter, the exercise can affect the decisions being made. Relatedly, the implicit assumptions about the type of data available may also be interpreted differently in a normative set-up: the descriptive interpretation these are assumptions about the availability of data in real life, while a normative reading can also apply to the data, calling for recording or reporting specific types of data.

Second, the revealed preference interpretation deals with the economist's problem, whereas the coherence test deals with the decision maker's problem. Put differently, the revealed preference question is meta-scientific: it isn't about which choices economic agents make, but which types of choices would make the economist conclude that a certain theory has been refuted. By contrast, the coherence test interpretation addresses the decisions that are being made by the decision maker herself. Being normative in nature, it attempts to capture the self-reflecting nature of a decision that needs to be justified to the self or to others.

Despite these differences, results proven in the revealed preference literature can be re-interpreted as coherence test results. At the same time, not all assumptions on models and observables are equally natural in both interpretations. The coherence test interpretation raises some questions which seem to us rather natural, and which have not received much attention in the revealed preference literature. We focus in this paper on settings in which an agent makes choices under uncertainty, where utilities are known or agreed-upon, but probabilities are not. Thus, the data are assumed to include real-valued acts defined on a given (and finite) set of states. Further, we assume that for each such act there exists a real-valued certainty equivalent, considered by the agent to be an evaluation of the act. The question is whether there exists an assignment of probabilities such that the valuations of the acts are consistent with either expected utility maximization or maxmin expected utility maximization.

The next subsection motivates the question we deal with, vis a vis better-studied questions in the revealed preference literature. Section 2 deals with the case of a Bayesian agent. Theorem 1 provides necessary and sufficient conditions under which there exists a probability vector such that the collection of acts and certainty equivalents is consistent with expected utility maximization relative to this probability. Theorem 2 in Section 3 establishes analogous conditions for the existence of a *set* of probability vectors, such that the collection of acts certainty equivalents is consistent with maxmin expected utility maximization. Section 4 places our work in the literature and Section 5 provides discussion.

1.2 Motivation

Given the focus of much of the revealed-preference literature on recovering *utilities* from observations of behavior, it may be a bit jarring to assume that utilities are known. When might utilities be observable? The coherence exercise may involve the balance sheet of an investment fund or financial institution, with allocations denoted in monetary values. Or, the institution may be a pension fund with risk guidelines that are implemented through the use of a specified CRRA utility function. The agents may be investors whose risk attitudes have been elicited by their investment counsellor. A public health official who has to make decisions about vaccination may be assumed to have a well-defined utility function (lives saved) but not necessarily a well-defined probability (over an epidemic outburst). Finally, a single individual evaluating her decision making may have insight into her utilities.

Revealed preference theory was developed in the context of certain outcomes, and so naturally focused on utilities. This focus carried over to the initial applications to settings of uncertainty, which assumed probabilities were given and again searched for utilities. By contrast, we are interested in situations where it is reasonable to assume that utilities are given but probabilities are not. We suspect that in many decision problems, the more daunting prospect is that of identifying probabilities. We expect the investment fund manager to be reasonably adept at identifying the possible outcomes of his ventures, but to find assigning probabilities to these outcomes more daunting. A political analyst may find it easier to assess the implications of various outcomes (such as a recession or a foreign policy breakthrough) for reelection than to attach probabilities to these outcomes.

The revealed preference literature typically assumes that the analyst observes a collection of feasible sets and the alternative chosen from each such feasible set. In our setting, information is instead conveyed by the specification of certainty equivalents. When might certainty equivalents be observable? The investment fund or financial institution may identify values of the various assets it holds. Alternatively, assets may be shifted between different units of a firm, fund or government on a break-even basis that implicitly reveals certainty equivalents. The agents may be laboratory subjects whose certainty equivalents may be elicited. A single individual evaluating her decision making may begin by identifying certainty equivalents. In a somewhat fanciful set-up, an official who makes decisions on behalf of others may be *required* to provide a protocol that suggests a decision theoretic model supporting her decision.

Even if observable, why reformulate the exercise in terms of certainty equivalents? The standard revealed preference application is described as allowing the analyst to observe a sequence of feasible sets and selected alternatives. Typically, however, the feasible set is *inferred* rather than observed. For example, a consumer who purchases consumption bundle x at price vector p is assumed to have had available any other bundle x' for which $px' \leq px$. In many applications, such as studying what goes into a grocery shopping cart, this seems eminently sensible. When examining choices under uncertainty, we regard this convention as leaning too heavily on the existence of complete markets. Consider a person wondering whether her decision to go to graduate school rather than pursue a career as a musician is consistent with her decision to cash out rather than invest an inheritance as well as her decision to move to the east coast. Or consider an investment fund manager faced with nonlinear prices (reflecting minimum position requirements, quantity discounts, liquidity restrictions, control considerations, and so on) for various assets. Either one may well find it impossible to fit these decisions into a traditional budget set. Instead, we find it more plausible that the agent begins her examination by converting the various options into a common currency of certainty equivalents.

2 Unambiguous Beliefs

In this section, we search for conditions under which there exists a probability vector for which a collection of decisions is consistent with expected utility maximization.

2.1 The Question

Let there be a set of states of the world, $S = \{1, ..., I\}$, indexed by i, and a set of acts $A = \{1, ..., K\}$, indexed by k. Act k is uncertain, yielding the utility x_i^k in state i, and thus corresponds to $x^k = \left(x_i^k\right)_{i \leq I}$. In each decision problem, the agent has to select an act out of a finite subset of acts. She is required to specify, for each act x^k, a certainty equivalent (CE) $c^k \in R$. It is implicitly assumed that in each decision problem the agent chooses an act k that maximizes c^k. Given a (finite) collection of act-CE pairs, $\left(x^k, c^k\right)$, when is there a probability over the states such that the expected x^k payoff, for each k, is c^k?

2.2 The Coherence Result

Appendix 6.1 proves the following:

Theorem 1 *Let there be given acts,* $x = \left(x^k\right)_{k \leq K}$ *(where* $x^k = \left(x_i^k\right)_{i \leq I} \in R^I$*) and a vector of certainty equivalents* $(c^k)_{k \leq K}$ *(where* $c^k \in R$*). The following are equivalent:*

(1.A) There exists a probability vector $p \in \Delta(S) \simeq \Delta^{I-1}$ *such that, for each*

$k \leq K$,[1]

$$p \cdot x^k = c^k.$$

(1.B) For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}$, *there exists a state* $i \leq I$ *such that*

$$\sum_{k=1}^{K} \lambda^k c^k \leq \sum_{k=1}^{K} \lambda^k x_i^k.$$

We mention that condition $(1.A)$ states a linear programming problem, which can be solved in polynomial time. Thus, for any given input (x, c), it is a task of manageable time complexity to determine whether condition $(1.A)$ holds and, if so, to find a probability vector that satisfies it. The theorem provides a general condition that characterizes those inputs for which the linear programming algorithm would yield a probability vector, versus those for which it would determine that no such vector exists.

2.3 Interpretation

The inequality in condition $(1.B)$ must hold for any collection of real numbers, $\left(\lambda^k\right)_{k \leq K}$, which may each be positive or negative. This means, in particular, that $(1.B)$ is equivalent to the reverse inequality

$(1.B')$ For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}$, there exists a state $i \leq I$ such that

$$\sum_{k=1}^{K} \lambda^k c^k \geq \sum_{k=1}^{K} \lambda^k x_i^k$$

and that $(1.B)$ and $(1.B')$ are each equivalent to the following condition,

$(1.B'')$ For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}$, there exist states $i, j \leq I$ such that

$$\sum_{k=1}^{K} \lambda^k x_j^k \leq \sum_{k=1}^{K} \lambda^k c^k \leq \sum_{k=1}^{K} \lambda^k x_i^k.$$

To see the meaning of condition $(1.B)$, first consider the weights $\lambda^k = 1$ for some k and $\lambda^l = 0$ for $l \neq k$. Then $(1.B'')$ boils down to

$$\min_i x_i^k \leq c^k \leq \max_i x_i^k, \tag{1}$$

which is obviously a necessary condition for c^k to be an expected value of x^k. That is, the certainly equivalent of the vector x^k cannot possibly be outside the range of the payoffs guaranteed by the vector.

[1] Here and in the sequel, the symbol $\cdot$ denotes an inner product.

Second, if there is only one decision to justify, that is, if $K = 1$, then (1) would also be sufficient. But the heart of the matter is that the decision maker needs to justify several decisions simultaneously. For example, suppose that there are two states ($I = 2$) and two acts ($K = 2$) given by the matrix

$$x = \begin{pmatrix} & i=1 & i=2 \\ k=1 & 1 & 0 \\ k=2 & 0 & 1 \end{pmatrix},$$

and the decision maker claims that the certainty equivalents of both x^1 and x^2 are $c^1 = c^2 = 0.4$. Obviously, each of these certainty equivalents can be justified on its own, but the two cannot be justified simultaneously. Condition ($1.B''$) would capture that by setting, say, $\lambda^1 = \lambda^2 = 0.5$ and observing that $\sum_{k=1}^{K} \lambda^k c^k = 0.4$ is outside of the range $\left[\min_i \sum_{k=1}^{K} \lambda^k x_i^k, \max_i \sum_{k=1}^{K} \lambda^k x_i^k\right] = [0.5, 0.5]$.

Third, to see that Condition ($1.B$) (or ($1.B'$) or ($1.B''$)) needs to resort to coefficients $\left(\lambda^k\right)_{k\leq K}$ that are not all of the same sign, consider the following example. Suppose that the decision maker attaches the certainty equivalent 0.5 to the vector $(1, 0)$ and attaches the certainty equivalent 0.8 to the vector $(2, -1)$. That is,[2]

$$x = \begin{pmatrix} & i=1 & i=2 \\ k=1 & 1 & 0 \\ k=2 & 2 & -1 \end{pmatrix} \sim c = \begin{pmatrix} k=1 & 0.5 \\ k=2 & 0.8 \end{pmatrix}.$$

Clearly, the first equivalence is only compatible with the probability vector $(0.5, 0.5)$ and the second is only equivalent with $(0.6, 0.4)$. Yet, if the $\left(\lambda^k\right)_{k\leq K}$ are all nonnegative, or all nonpositive, this assignment of certainty equivalents will not be ruled out by condition ($1.B$). To see this, suppose that the $\left(\lambda^k\right)_{k\leq K}$ are all nonnegative, and are not all 0 (in which case the equality in condition ($1.B$) obviously holds). Then there is no loss of generality in assuming $\lambda^1 + \lambda^2 = 1$, and in this case we have

$$\left[\min_i \sum_{k=1}^{K} \lambda^k x_i^k, \max_i \sum_{k=1}^{K} \lambda^k x_i^k\right] \supseteq [0, 1]$$

and $\sum_{k=1}^{K} \lambda^k c^k \in [0.5, 0.8]$ lies within this range, thus satisfying condition ($1.B$). By contrast, choosing $\lambda^1 = 3, \lambda^2 = -1$ we find that

$$\left[\min_i \sum_{k=1}^{K} \lambda^k x_i^k, \max_i \sum_{k=1}^{K} \lambda^k x_i^k\right] = \{1\},$$

and $\sum_{k=1}^{K} \lambda^k c^k = 0.7$ falls outside this singleton range, violating condition ($1.B$).

[2]The equivalence sign between a matrix and a vector should be read to hold row-by-row.

3 Ambiguous Beliefs

Much of the motivation for our inquiry arises out of the belief that people may more readily assess utilities than probabilities. For the same reasons that probabilities are elusive, a person may be reluctant to commit to a single probability vector, being instead more comfortable thinking in terms of a set of probability measures. We therefore turn to search for conditions under which there exists a *set* of probability vectors under which a collection of decisions are consistent with maxmin expected utility maximization.

Suppose, for example, an investment fund manager is called upon to defend her decisions. She might respond,

> "There was a lot about the markets that I did not understand, and I found it best neither to buy nor to sell. I know the collection of the decisions I made cannot be justified by a single probability, but I didn't have a single probability that I could trust. Indeed, were I to invest your money based on one such probability you would have asked me where I got it from. Instead, I used a model that allows for ambiguity, and, with some ambiguity aversion, I can explain why I did what I did."

Some of the investor's clients might have wished to know a priori that she was prone to making decisions by non-Bayesian models, and some might have preferred to work with a Bayesian fund manager. Other investors may be skeptical of a manager who claims they can identify probabilities precisely, and may prefer a manager who is clear about the relevant ambiguity. Similarly, an individual may be more comfortable assessing a sequence of decisions by appealing to bounds on probabilities but not by making point estimates of probabilities.

3.1 The Question

The setting matches that of Section 2.1. We have a set of states of the world, $S = \{1, ..., I\}$, indexed by i, and a set of acts $A = \{1, ..., K\}$, indexed by k. Each act is a utility vector $x^k = \left(x_i^k\right)_{i \leq I}$ and there is a certainty equivalent c^k attached to it. As above, we implicitly assume that the act chosen is a maximizer of the certainty equivalent.

We now allow beliefs to incorporate ambiguity. Hence, we ask whether there exists a *set* of probability vectors such that the putative certainty equivalent for each act is the minimum (over the set of probabilities) expected utility of that act. More explicitly, we ask, given $x = \left(x^k\right)_{k \leq K}$ (where $x^k = \left(x_i^k\right)_{i \leq I} \in R^I$) and a vector of certainty equivalents $(c^k)_{k \leq K}$ (where $c^k \in R$) whether there exists a convex, closed set of probabilities on S, $C \subset \Delta(S) \simeq \Delta^{I-1}$ such that, for each $k \leq K$,

$$\min_{p \in C} p \cdot x^k = c^k$$

where $p \cdot x^k = \sum_{i \leq I} p_i x_i^k$.

3.2 The Coherence Result

Appendix 6.2 proves the following:

Theorem 2 *Let there be given acts, $x = \left(x^k\right)_{k \leq K}$ (where $x^k = \left(x_i^k\right)_{i \leq I} \in R^I$) and a vector of certainty equivalents $(c^k)_{k \leq K}$ (where $c^k \in R$). The following are equivalent:*

(2.A) *There exists a convex, closed set of probabilities on S, $C \subset \Delta(S) \simeq \Delta^{I-1}$ such that, for each $k \leq K$,*

$$\min_{p \in C} p \cdot x^k = c^k.$$

(2.B) *(i) For every collection of non-negative real numbers $\left(\lambda^k\right)_{k \leq K}$, there exists a selection of a state $i \leq I$ such that*

$$\sum_{k=1}^{K} c^k \lambda^k \leq \sum_{k=1}^{K} x_i^k \lambda^k.$$

(ii) For every act $l \leq K$, and collection of non-negative real numbers $\left(\lambda^k\right)_{k \leq K, k \neq l}$, there exists a state $i = i_l \leq I$, such that

$$\sum_{\substack{k \neq l \\ k \leq K}} \lambda^k \left(c^k - x_{i_l}^k\right) \leq c^l - x_{i_l}^l.$$

As in the previous result, $(2.A)$ states a linear programming problem, and can therefore be verified in polynomial time. Indeed, if the condition holds, the relevant probabilities can also be found in polynomial time. Finally, as in the Bayesian case, the general characterization theorem allows us to determine for which inputs (x, c) the linear programming algorithm will answer the existence question in the affirmative.

3.3 Interpretation

Condition 2.$B(i)$ is a sort of feasibility condition, guaranteeing that the certainty equivalents c^k are not too high relative to the given x's. It says that, for every weighted "utilitarian" average, that is, every "average" act, what is required (on the left side) does not exceed what this "mixed act" yields *in at least one state*. This is somewhat reminiscent of the first characterization result (Theorem 1) in Gilboa, Samuelson and Schmeidler (2014).

Condition 2.$B(ii)$ has to do with comparisons among the acts. It says that, for every act l and for every average of the *other* acts, there is at least one state i_l in which the gain $c^l - x^l_{i_l}$ provided by l's certainty equivalent c^l must be at least as large as the gain received by the average of the other acts. In other words, it cannot be the case that in each and every state, moving to the certainty equivalent is worse for act l than it is for the average of the other acts.

To see why we might expect this second condition to hold, suppose there just two acts, l and k, and two states, 1 and 2. In this case, acts k plays the role of the average of the acts other than l. Suppose there exists a set C, in this case given by an interval of the form $[\underline{p}_1, \overline{p}_1]$ of possible probabilities of state 1, such that condition 2.A holds. Suppose $x^k_1 < x^k_2$. Then act k is evaluated according to the probability $\overline{p}_1$, since k yields higher payoff in state 2 than in state 1, and thus the minimal expected payoff is obtained when the probability to state 1 is the largest. Hence, the certainty equivalent c^k satisfies

$$\overline{p}_1(c^k - x^k_1) + (1 - \overline{p}_1)(c^k - x^2_k) = 0,$$

where

$$\begin{aligned} c^k - x^k_1 &> 0 \\ c^k - x^k_2 &< 0. \end{aligned}$$

Now we consider two possibilities. First, it may be that we also have $x^l_1 < x^l_2$. Then act l is also evaluated with the probability $\overline{p}_1$, since l yields a higher payoff in state 2 and hence obtains the minimal expectation when the probability of state 1 is the highest. Hence, the certainty equivalent c^l also satisfies

$$\overline{p}_1(c^l - x^l_1) + (1 - \overline{p}_1)(c^l - x^l_2) = 0.$$

But then we have two identical linear equations describing the change from the acts in x to their certainty equivalents, and hence it is impossible that

$$\begin{aligned} c^k - x^k_1 &> c^l - x^l_1 \\ c^k - x^k_2 &> c^l - x^l_2, \end{aligned}$$

which is what is prohibited by condition $2.B(ii)$. When acts k and l are evaluated by the same probabilities, it cannot be the case that more is gained *in every state* by moving from act k (rather than l) to its certainty equivalent.

Alternatively, it may be that we have $x_1^l > x_2^l$. Then act l is evaluated by probability $\underline{p}_1$, since l yields a higher payoff in state 1 and obtains the minimum expectation when this state's probability is minimal. Hence, the certainty equivalent c^l also satisfies

$$\underline{p}_1(c^l - x_1^l) + (1 - \underline{p}_1)(c^l - x_2^l) = 0,$$

where

$$\begin{aligned} c_1^l - x_1^l &< 0 \\ c_2^l - x_2^l &> 0. \end{aligned}$$

But then we have $c_2^k - x_2^k < 0 < c_2^l - x_2^l$, as required by condition $2.B(ii)$. Our condition $2.B(ii)$ is the generalization of this observation to more than two acts and two states.

Suppose we have acts A and B and an arbitrary number of states. We show that if

$$\min_{p \in C} p^A x^A = c^A, \qquad \min_{p \in C} p^B x^B = c^B$$

then for any $\lambda^B > 0$, there exists a state i with

$$\lambda^B(c^B - x_i^B) \leq c^A - x_i^A, \tag{2}$$

as required by condition $(2.B(ii))$. Notice first that it suffices to show the result for $\lambda^B = 1$.[3] Next, it sacrifices no generality to assume that $c^A = c^B$, because the objective functions are linear and we can subtract $c^B - c^A$ from every one of B's payoffs to set their certainty equivalents equal. Let c denote the common value $c^A = c^B$.

We proceed by contradiction. Suppose we have $p^B x^B = c$, but there is no state i satisfying (2), i.e., for every state i, we have

$$c - x_i^B > c - x_i^A.$$

Then for all i, we have $x_i^A > x_i^B$, and hence

$$p^A x^A > p^A x^B \geq \min_{p \in C} p x^B = p^B x^B = c,$$

[3]Faced with any other value of λ^B, first divide all A's payoffs by λ^B, then establish the result with the modified payoffs for $\lambda^B = 1$, at which point multiplying A's payoffs by λ^B gives the result for the value of λ^B. Hence, we set $\lambda^B = 1$.

which is a contradiction.

By itself, this argument does not provide particularly compelling evidence for condition $2.B(ii)$. In particular, in the two-act case, it is clear that there must also exist a state j such that

$$\lambda^B(c^B - x_J^B) \geq c^A - x_J^A.$$

To see this, we need only reverse the roles of acts A and B in the preceding argument. How do we know we have chosen the right property from the two-act case to generalize? How do we know that a counterpart of condition $2.B(ii)$ does not also hold, with the inequality reversed?

We show by example that with three acts and two states, it can be the case that for every state

$$\sum_{\substack{k \neq l \\ k \leq K}} \lambda^k \left(c^k - x_{i_l}^k\right) < c^l - x_{i_l}^l.$$

Hence, we cannot establish a counterpart of condition $2.B(ii)$ with the inequality reversed. The allocation in the example is given by

$$\begin{array}{ccc} & i=1 & i=2 \\ k=A & 1 & 3 \\ k=B & 3 & 1 \\ k=C & 2 & 2 \end{array}.$$

The set of beliefs C is given by the entire unit interval. Hence, we have

$$c^A = 1, \qquad c^B = 1, \qquad c^C = 2.$$

Now for state 1 (state 2 is analogous) we have

$$\begin{aligned} \frac{1}{2}(c^A - x_1^A) + \frac{1}{2}(c^B - x_1^B) &= \frac{1}{2}(0) + \frac{1}{2}(-2) \\ &= -1 \\ &< 0 \\ &= c^C - x_1^C, \end{aligned}$$

giving the result.

4 Relationship to the Literature

In this section we fit our analysis into the revealed-preference literature. We will frequently refer to Chambers and Echenique (2016), whose presence makes our task much easier and to which we often refer, with our apologies to those whose original papers we would have otherwise cited.

To organize ideas, we note that a revealed preference analysis consists of:

- A set of objects, such a set of commodity vectors or a set of lotteries or a set of state-contingent consumption plans;
- A (typically finite) set of data, such as a set of feasible sets of objects and the designation of the object chosen from each set;
- A question, such as whether the data are consistent with the maximization of a utility function; and
- A result, typically in the form of a representation theorem, identifying properties of the data that are necessary and sufficient for the question to have an affirmative answer.

The first generation of revealed-preference papers, pioneered by Samuelson (1938) and Houthakker (1950) and examined in the first six chapters of Chambers and Echenique (2016), addresses cases in which there is no uncertainty. The data typically consist of a finite collection of consumer budget sets and an indication of which consumption bundle was chosen out of each set, with the question being whether these choices are consistent with a preference ordering or the maximization of a utility function.

A next generation of papers examines settings in which an agent is choosing under uncertainty, but objective probabilities are given. The data typically consist of a collection of feasible sets of lotteries or state-contingent consumption plans and an indication of the choice from each feasible set. The question is again whether these choices are consistent with a preference ordering or the maximization of (for example) an expected utility function, given the specified probabilities. Figure 1 summarizes the central papers. [4]

The natural next step is to assume that neither probabilities nor utilities are given. Figure 2 summarizes these papers. We discuss the three most relevant papers in somewhat more detail.

The fundamental result here is by Echenique and Saito (2015) (see Chambers and Echenique (2016, Section 8.2) for an exposition). Echenique and Saito as-

[4] Within this setting, Polisson, Quah and Renou (2019) address a somewhat more general methodological question. They distinguish between two revealed-preference settings. In the finite case, the data consist of a finite set $\{x_k, y_k\}_k$ of pairs of allocations with the property that $x_k \succsim y_k$, and another finite set $\{x_k, y_k\}_k$ of pairs of allocations with the property that $x_k \succ y_k$. In the infinite case the data consist of a set $\{x_t, B_t\}_t$, where each B_t is a (typically infinite) budget set and x_t is the allocation chosen from this budget set. In principle, the analysis is easier in the finite case, since rationalizing the data then involves finding a solution to a finite number of inequalities, while infinite budget sets give rise to an infinite collection of inequalities. The central result in Polisson, Quah and Renou (Theorem 1) shows that for any infinite problem $\{x_t, B_t\}_t$, one can find a finite set of inequalities such that there exists a belief and utility function rationalizing choices in the infinite problem if and only if they satisfy the finite set of inequalities. This effectively reduces the complexity of the infinite problem to that of the finite problem.

Source	Objects	Data	Question	Result
Chambers and Echenique (2016, Section 8.1.2); Fishburn (1975)	Lotteries over arbitrary prize set; objective probabilities	List of choices from pairs of lotteries	Does there exist an expected utility representation?	Yes, if there exists no mixture yielding two identical compound lotteries, with the constituents of one mixture pairwise preferred to those of the other.
Chambers and Echenique (2016, Section 8.1.3); Kubler, Selden and Wei (2014)	State contingent monetary payoffs; objective probabilities	List of budget sets and selected plans	Does there exist a risk-averse expected utility representation?	Yes, if products of ratios of risk-neutral prices exhibit no cycles.
Chambers, Liu and Martinez (2014)	State contingent consumption plans, allocations drawn from Euclidean space; objective probabilities	List of budget sets and selected plans	Does there exist a risk-averse expected utility representation?	Yes, if there exists no mixture yielding two identical compound lotteries, with the constituents of one mixture pairwise preferred to those of the other.
Kubler, Selden and Wei (2014)	State contingent monetary payoffs; objective probabilities	Demand functions, giving contingent plan as function of prices, probabilities and income	Does there exist an expected utility representation?	Yes, if demands are functions of ratios of probability-weighted prices
Lin (2019)	Lotteries over arbitrary prize set; objective probabilities	List of finite feasible sets and selected lottery	Does there exist a betweenness representation?	Yes, if a geometric condition is satisfied.
Polisson, Quah and Renou (2019, Section 2)	State contingent consumption plans, allocations drawn from Euclidean space, objective probabilities	List of feasible (e.g., need not be linear) sets and selected plans	Does there exist an expected utility representation?	Necessary and sufficient conditions can be expressed in terms of a finite lattice of consumption plans

Figure 1: Summary of reveled-preference papers examining uncertainty with objective probabilities.

Paper	Objects	Data	Question	Answer
Chambers and Echinique (2016, Section 8.2), based on Echinique and Saito (2015)	State contingent monetary payoffs	List of budget sets and selected plans	Does there exist a risk averse subjective expected utility consistent with data?	Yes, if products of ratios of prices exhibit no cycles.
Chambers and Echinique (2016, Section 8.3); Börgers (1993)	State contingent plans, abstract prize set	Preference relation $\succsim$ over subset F of X^Ω	If $\succsim$ identifies unique best element of F, does there exist subjective expected utility representaiton that uniquely picks this best element?	Yes, if $\succsim$ monotonic.
Chambers and Echinique (2016, Section 8.4), drawing on Bossert and Suzumura (2012)	State contingent plans, abstract prize set	Preference relation $\succsim$ over finite X^Ω	Does there exist act-dependent-probability subjective expected utility representation of $\succsim$? (Allows probabilities to depend on acts in arbitrary ways.)	Yes, if $\succsim$ satisfies the weak condition of uniform monotonicity.
Polisson, Quah and Renou (2019, Section A2.2)	State contingent consumption plans, allocations drawn from Euclidean space	List of feasible (not necessarily linear) sets and selected plans	Does there exist a maxmin subjective expected utility representation?	Yes, if our Condition (2) in the proof of Theorem 2 is satisfied.

Figure 2: Summary of reveled-preference papers examining uncertainty without objective probabilities.

sume that x^k takes on monetary values, and offer a characterization of expected utility maximization that is similar to the result in Chambers and Echenique (2016, Section 8.1.3) (for objective probabilities). In particular, the necessary and sufficient condition is a restriction on products of ratios of risk-neutral prices. The risk-neutral prices under uncertainty play the role otherwise played by objective probabilities. This derivation and result depends importantly on the assumption that feasible sets are budget sets. In contrast, we follow Polisson, Quah and Renou (2019) in assuming that the data consist not of alternatives chosen from budget sets, but rather of pairs of (weakly or strictly) ranked alternatives. (We in addition assume that one such alternative in each pair is a certainty equivalent.) As we explain in Section 1, budget sets are typically at best inferred, and are unlikely to be immediately apparent once we push the domain of revealed preference theory outside the realm of traditional consumer theory. Instead, we suspect normative applications of revealed preference theory are more likely to confront a collection of allocations, and (if all goes well) their associated certainty equivalents.

Chambers and Echenique (2016, Section 8.4) examine the rationalization of choice between acts by formulations in which the probability with which the expected utility of an act is evaluated can depend on the act. Notice that maxmin utility is one such formulation, as the probability p by which an act x is evaluated is given by $\arg\min_{p \in C} px$ for some set C. In general, probabilities can depend on acts in many different ways. The key result is their Proposition 8.9, stating that a collection of reveled preferences $\succsim_R$ has an expected utility representation in which probabilities can depend on acts if and only if the $\succsim_R$ satisfies a quite weak condition known as uniform monotonicity.[5] The notion of uniform monotonicity is characterized in Bossert and Suzumura (2012).[6] The proof relies critically on the ability to assign to any act x an arbitrary probability distribution by which it is to be evaluated. In contrast, maxmin expected utility puts restrictions on which probabilities are used to evaluate which acts.

In Section A2.2 on their online appendix, Polisson, Quah and Renou (2019) examine maxmin utility. Their Proposition A.1 shows that there exists a maxmin utility representation for a set of data $\{x_t, B_t\}_t$ if and only if there is a set of probabilities and a finite collection of utilities such that these probabilities and utilities solve inequalities (their A.3–A.5) that are the counterpart of our condition (2) in the proof of Theorem 2. Polisson, Quah and Renou do not develop conditions for Condition 2 to be satisfied (as we do), while arguing that Condition 2 can be applied to the infinite case (while we consider only the finite case).

Finally, two papers adopt our perspective of assuming that utilities are given and focussing on the existence of probabilities.

Chambers, Echenique and Saito (2016) assume that there is uncertainty, no probabilities are given, and the data consist of a collection $\{x^k, p^k\}_k$, where p^k is a price vector and the presumption is that the state-contingent plan x^k was

[5] Choices satisfy uniform monotonicity if, for two acts x and y, if $x(\omega) \succsim_R y(\omega')$ for all states ω and ω' (think of this as the indication that the worst outcome under x, taken as a constant act, is better than the best outcome under y, also taken as a constant act), then $x \succsim_R y$.

[6] Bossert and Suzumura (2012) assume there is uncertainty, no probabilities are given, and choice sets are arbitrary sets of state-contingent consumption plans, taking either monetary values or identifying commodity allocations. They are concerned with when there exists a preference relation exhibiting certain general properties that rationalize the data, but do not investigate whether these preferences might correspond to something like expected utility maximization or maxmin expected utility maximization.

chosen from the set $\{x : p^k x \leq p^k x^k\}$. The vectors x^k either take on utility values, or monetary values for a risk-neutral agent. Chambers, Echenique and Saito (2016) offer necessary and sufficient conditions (Theorem 3) for the data to be characterized by maxmin utility. The sufficiency of the conditions they offer is almost immediate, in that it allows one to directly reinterpret the prices attached to the various states as probabilities.

Arieli and Mueller-Frank (2017) posit the existence of compact metric spaces Ω of states and A of actions, and the existence of a (known) utility function $u : \Omega \times A \to \mathbb{R}$. They ask whether, knowing the utility function u and having observed the optimal action a, one can infer the beliefs of the agent. The answer is affirmative, for generic utility functions u, if the set A contains no isolated points. This result reflects the fact that knowing the function u on a rich domain (ensured by the absence of isolated points) provides a great deal of information.

The present paper is related to Gilboa and Samuelson (2020), which extends the analysis of Gilboa, Samuelson, and Schmeidler (2014) to ambiguous beliefs. The latter considered trades among Bayesian agents who hold different beliefs, and suggested to distinguish between those trades that could be simultaneously justified, for all agents involved, by a shared belief and those that could not, that is, that crucially hinge upon differences in opinions. Gilboa and Samuelson (2020) considers the same notion of "no-betting" Pareto domination, but allows the agents to hold non-Bayesian beliefs. Specifically, it assumes that the agents maximize maxmin expected utility, and asks when a proposed trade can be simultaneously justified, for all agents involved, by a shared set of probabilities. Hence it shares with the present paper a normative approach, dealing with collections of decisions that can or cannot be justified. However, the two have rather different motivations, which lead to important differences in the models. Specifically, Gilboa and Samuelson's (2020) focus is naturally-occurring trade, and it thus does not assume that pre- and post-trade bundles be declared equivalent to risk-free ones. By, contrast, in this paper our focus is a single decision maker, and we assume we may demand that they evaluate each possible act and document its expected utility. This implies that the two papers have rather different assumptions regarding observable data, and this, in turn, renders their results independent.

5 Discussion

A textbook application of normative decision theory suggests that the decision maker specify her problem—including decision variables, objective functions, constraints, and beliefs—and invoke an algorithm to solve it. In such a scenario there is no need to ask whether one can find parameters, such as beliefs, that justify a given decision. There is no "given decision", with the decision maker instead solving for the best decision. This scenario characterizes many decisions, which are sometimes so routine as to be taken for granted. For example, suppose Mary needs to drive from city A to city B. The set of available options is clearly defined by a graph. The uncertainty involved may well be amenable to statistical analysis – Mary can collect data not only about length of paths but also about congestion, leading to rather accurate estimates of probability distributions. Coupling this with some reflection as to her objective, Mary can perform a reasonably straightforward analysis. Such an application is a triumph of normative decision theory.

In many other scenarios, there is no obviously "correct" model of the problem. Suppose that John is about to graduate from high school and has to choose a career path. Identifying the set of options is a daunting task in itself. John might apply to various colleges, choose from a variety of possible majors, contemplate various professions, choose to attend or forego a graduate school, and so on. John may also consider a professional career as a tennis player, in combination with or instead of college. Having formed the list, it is not at all obvious how John is to assign probabilities to the resulting outcomes. Moreover, the decision problem is also replete with uncertainty about the external world that is not easily quantifiable by probabilities. What will be the effect of autonomous vehicles on various career choices that John can make? Or the impact of global warming?

It might seem that decision theory can help in well-defined problems such as Mary's, but that it is useless in more vague problems such as John's. This paper views such a conclusion as premature (as do Gilboa, Postlewaite, Samuelson, and Schmeidler (2018) and Gilboa, Rouziou, and Sibony (2018)). Decision theory can help with even the vaguest of problems by acting as a *coherence test*. We can imagine John making up his mind by whatever mix of intuition, advice, imitation, calculation, and guesswork he has available, and then checking his judgement by asking whether he can justify the decision using a model. In this paper we assume that the set of acts, the utility function, the constraints and the states of the world are relatively straightforward. We ask whether, given the other parameters, there exist beliefs that justify John's tentative choice.

Conceptually, we imagine a dialog between the decision making agent and decision theory. In scenarios such as Mary's, the dialog is rather simple: the decision maker provides a decision theorist (or a piece of software) the relevant parameters, and gets back an optimal act. Such a normative application may be so successful as to become a smartphone application (no pun intended): it can be fully automated and relegated to machines. In other cases, such as John's, the decision maker may offer a potential alternative and ask to see what must be assumed in order to justify it. In between these extremes, there may be room for a dialog in which an attempt to justify a tentative decision helps to clarify the problems and perhaps lead to another decision.

The need to justify a decision by a formal model may therefore help individuals reach decisions that they end up liking better than those suggested by their intuition. More importantly, when we consider agents who are part of institutions, the need for justification goes beyond the agents' own peace of mind; it may be an essential part of responsible, transparent decision making. As suggested in the Section 1, we can consider an agent whose job is to invest other people's money. Suppose Sarah is a hedge fund manager, making financial trading decisions on behalf of her clients. The available options may be well defined and the utility function of her clients may be assumed to be monotone in wealth, with risk attitudes that Sarah may have elicited via questionnaires, interviews, and experience. But Sarah deals with uncertainty that is hard to quantify. Even when considering investment horizon of a few months or years, there are many uncertainties to deal with, and these are not as easily quantifiable as is driving time in Mary's problem.

Indeed, if all uncertainties about financial markets were amenable to statistical analysis, the hedge fund management problem could have been relegated to a smartphone app as well. If Sarah makes a living as a hedge fund manager, we may assume that some people trust her intuition and her ability to "read" the markets. At the same time, Sarah's clients would not be happy with just any decision she makes. If they end up losing money, they may go back to her investment decisions, and ask, "But why did you do that? *What were you thinking* exactly?"

The present paper provides the tools for Sarah to determine which decisions she could justify if called upon to do so. While the question is formally a revealed preference type of problem, it has a different interpretation: rather than asking a methodological question, along the lines of "which data sets would make an economist admit that a theory has been refuted?" we ask, "which data sets could be justified by the agent?" This different interpretation leads to a change of focus, and leads us to problems in which it seems natural to assume that utility is given, and that the focus of analysis is probabilities.

We believe that in many cases there are obvious benefits to protocols that require justification of decisions by decision-theoretic models. Before a pension fund invests in exotic assets, before a president decides to bomb or invade another country, before countries embrace or abandon climate change policies, we would ask that they put the tools provided in this paper to work and identify at least one set of coherent probabilities rationalizing their decision. Clearly, one can also point to drawbacks of such protocols, ranging from cumbersome bureaucracy to corruption. We view our contribution as raising the question, and laying some theoretical foundations for decision protocols that might be needed should one choose to implement them.

6 Appendix: Proofs

6.1 Proof of Theorem 1

The proof presents seven conditions, the first of which is statement $(1.A)$ of Theorem 1 and the last of which is statement $(1.B)$ of Theorem 1, and shows that each successive pair of conditions is equivalent.

- **Condition 1:** There exists a probability $p \in \Delta(S)$ such that, for each $k \leq m$,
$$p \cdot x^k = c^k.$$

- **Condition 2:** The following linear programming problem is feasible:
$$\begin{aligned} (P) \qquad & \min_{(p_i)_{i \leq I}} \sum_{i=1}^{I} 0 \cdot p_i \\ s.t. \qquad & \sum_{i=1}^{I} x_i^k \cdot p_i = c^k \qquad \forall k \leq K \\ & \sum_{i=1}^{I} p_i = 1 \\ & p_i \geq 0 \qquad \forall i \leq I. \end{aligned}$$

Condition (1) and Condition (2) are equivalent because the feasible set in (P) is simply the statement that (p_i) is a probability vector, according to which the expectation of each $\left(x_i^k\right)_{i \leq I}$ is c^k.

- **Condition 3:** The following LP problem is bounded:
$$\begin{aligned} (D) \qquad & Max_{(\lambda^k)_k, \mu} \left[\left(\sum_{k=1}^{K} c^k \lambda^k \right) + \mu \right] \\ & s.t. \sum_{k=1}^{K} x_i^k \lambda^k + \mu \leq 0 \qquad \forall i \leq I \end{aligned}$$

  (with $\left(\lambda^k\right)_k, \mu$ unconstrained).

Condition (2) and Condition (3) are equivalent because (D) is the dual of (P).

- **Condition 4:** The LP problem (D) in Condition 3 is bounded above by zero.

Condition (3) and Condition (4) are equivalent because the feasible set of (D) is a cone.

- **Condition 5:** For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}, \mu$,

$$\begin{array}{lll} \text{IF} & \mu \leq -\sum_{k=1}^{K} x_i^k \lambda^k & \forall i \leq n \\ \text{THEN} & \left(\sum_{k=1}^{K} c^k \lambda^k\right) + \mu \leq 0. & \end{array}$$

Condition (4) and Condition (5) are equivalent because (5) simply says that each point $\left(\left(\lambda^k\right)_{k \leq K}, \mu\right)$ which is feasible for (D) obtains a non-positive objective function value in it.

- **Condition 6:** For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}$,

$$\sum_{k=1}^{K} c^k \lambda^k \leq \max_{i \leq n} \sum_{k=1}^{K} x_i^k \lambda^k.$$

Condition (6) is equivalent to Condition (5) because in (5) we might set μ to its maximal value (for any $\left(\lambda^k\right)_{k \leq K}$), that is $\min_i \left(-\sum_{k=1}^{K} x_i^k \lambda^k\right)$, and then require that $\left(\sum_{k=1}^{K} c^k \lambda^k\right) + \min_i \left(-\sum_{k=1}^{K} x_i^k \lambda^k\right) \leq 0$.

- **Condition 7:** For every collection of real numbers $\left(\lambda^k\right)_{k \leq K}$, there exists a state $i \leq n$ such that

$$\sum_{k=1}^{K} c^k \lambda^k \leq \sum_{k=1}^{K} x_i^k \lambda^k$$

Condition (7) is simply a re-statement of Condition (6).

6.2 Proof of Theorem 2

The proof presents ten conditions, the first of which is statement $(2.A)$ of Theorem 2 and the last of which is statement $(2.B)$ of Theorem 2, and shows that each successive pair of conditions is equivalent.

- **Condition 1:** There exists a convex, closed set of probabilities on S, $C \subset \Delta(S)$ such that, for each $k \leq m$,

$$\min_{p \in C} p \cdot x^k = c^k.$$

- **Condition 2:** For every $k \leq K$ there exists $p^k \in \Delta(S)$ such that,

$$p^k \cdot x^k = c^k$$

and, for each $l \leq K$,

$$p^l \cdot x^k \geq c^k.$$

Proof that Conditions 1 and 2 are equivalent: If Condition 1 holds, we can pick, for each k, a probability p^k in C that achieves the minimum expected payoff for k, and these probabilities satisfy Condition 2. Conversely, if Condition 2 holds, $C = conv\left(\left\{p^k\right\}_{k \leq K}\right)$ satisfies Condition 1.

- **Condition 3:** The following linear programming problem is feasible:

$$(P) \qquad Min_{\left(p_i^k\right)_{i \leq I, k \leq K}} \sum_{k=1}^{K} \sum_{i=1}^{I} 0 \cdot p_i^k$$

$$s.t.$$

$$\begin{aligned}
\sum_{i=1}^{I} x_i^k \cdot p_i^k &= c^k \qquad \forall k \leq K \\
\sum_{i=1}^{I} x_i^k \cdot p_i^l &\geq c^k \qquad \forall k, l \leq K, \quad k \neq l \\
\sum_{i=1}^{I} p_i^k &= 1 \qquad \forall k \leq K \\
p_i^k &\geq 0 \qquad \forall i \leq I, k \leq K.
\end{aligned}$$

Proof that Conditions 2 and 3 are equivalent: The feasible set of Problem (P) in Condition 3 is a restatement of Condition 2. Hence, Problem (P) is feasible if and only if Condition 2 is satisfied. The objective in Problem (P) plays no role in this equivalence; we select a zero objective function for convenience.

- **Condition 4:** The following linear programming problem is bounded:

$$
\begin{aligned}
&(D) \qquad Max_{(\lambda^{kl})_{k,l\leq K},(\mu^k)_{k\leq K}} \left[\left(\sum_{k=1}^{K}\sum_{l=1}^{K} c^k\lambda^{kl}\right)+\sum_{k=1}^{K}\mu^k\right] \\
&s.t. \\
\sum_{l=1}^{K} x_i^l \lambda^{lk} + \mu^k &\leq 0 \qquad \forall i \leq I, k \leq K \\
\lambda^{kl} &\geq 0 \qquad \forall k,l \leq K, \quad k \neq l.
\end{aligned}
$$

Proof that Conditions (3) and (4) are equivalent: One has to verify that Problem (D) in Condition 4 is indeed the dual problem of Problem (P) in Condition 3, and then employ the weak duality theorem to conclude that the primal is feasible if and only if the dual is bounded.

To see that Problem (D) is the dual of Problem (P), observe the following:

- Problem (P) has KI variables, $\left(p_i^k\right)_{i\leq I,k\leq K}$.

  Problem (P) has K^2+K constraints (apart from the non-negativity constraints): K^2 constraints of the form $\sum_{i=1}^{I} x_i^k \cdot p_i^l \geq c^k$, with the weak inequality (the natural inequality for a minimization problem) replaced by an equality on the diagonal; and I equality constraints ensuring that the p^k are probability vectors.

- Each of the first K^2 constraints is associated with a variable λ^{kl}, with the diagonal variables unconstrained (corresponding to equalities in Problem (P)), and the others are non-negative (corresponding to natural $\geq$ constraints in Problem (P)). Hence the variable λ^{kl} is attached to $\sum_{i=1}^{I} x_i^k \cdot p_i^l \geq c^k$ for $k \neq l$, and the variable λ^{kk} is attached to $\sum_{i=1}^{I} x_i^k \cdot p_i^k = c^k$. Each of the last K constraints is associated with a variable μ^k which is unconstrained, as the corresponding constraint in Problem (P) is an equality.

  Problem (D) is a maximization problem with an objective function whose coefficients are the right side of the constraints in Problem (P): for each k, all of variables $\left(\lambda^{kl}\right)_l$ have coefficient c^k, and each μ^k is multiplied by 1.

  The right side of all constraints in Problem (D) is zero, following our choice of the objective function coefficients in (P).

- A tedious but straightforward accounting of the subscripts and superscripts confirms that the constraint coefficient matrix in Problem (D) is the transpose of that of Problem (P).

- **Condition 5:** The linear programming problem (D) in Condition 4 is bounded above by zero.

Proof that Conditions (4) and (5) are equivalent: The feasible set of Problem (D) is a cone, so that $\left(\left(\lambda^{kl}\right)_{k,l\leq K},\left(\mu^{k}\right)_{k\leq K}\right)$ is feasible if and only if any positive multiple of it is. Thus, if there exists a point $\left(\left(\lambda^{kl}\right)_{k,l\leq K},\left(\mu^{k}\right)_{k\leq K}\right)$ that obtains a positive value of the objective function, (D) is unbounded, contradicting Condition (4). Otherwise, (D) is bounded by zero.

- **Condition 6:** For every collection of real numbers $\left(\lambda^{kl}\right)_{k,l\leq K},\left(\mu^{k}\right)_{k\leq K}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$,

$$\text{IF} \qquad \mu^k \leq -\sum_{l=1}^{K} x_i^l \lambda^{lk} \qquad \forall i \leq I, k \leq K$$

$$\text{THEN} \qquad \left(\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl}\right) + \sum_{k=1}^{K} \mu^k \leq 0.$$

Proof that Conditions (5) and (6) are equivalent: The feasible set of Problem (D) contains all the matrices $\left(\lambda^{kl}\right)_{k,l\leq m}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$, and the vectors $\left(\mu^{k}\right)_{k\leq m}$ such that

$$\mu^k \leq -\sum_{l=1}^{K} x_i^l \lambda^{lk} \qquad \forall i \leq I \tag{3}$$

and, conversely, any such matrix and vector defines a point in the feasible set of Problem (D). The consequent of Condition (6) then states that the objective function isn't positive for such a point.

- **Condition 7:** For every collection of real numbers $\left(\lambda^{kl}\right)_{k,l\leq K}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$,

$$\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl} \leq \sum_{k=1}^{K}\left(\max_{i\leq n}\sum_{r=1}^{K} x_i^r \lambda^{rk}\right). \tag{4}$$

Proof that Conditions (6) and (7) are equivalent: Condition (6) is equivalent to the statement that

$$\left(\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl}\right) + \sum_{k=1}^{K} \mu^k \leq 0$$

holds for the maximal values of $\left(\mu^k\right)_{k\leq K}$ permitted by the antecedent of Condition (6), which, by (3), are

$$\min_{i\leq I}\left(-\sum_{l=1}^{K} x_i^l \lambda^{lk}\right).$$

It thus suffices to show that (4) is equivalent to

$$\left(\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl}\right) + \sum_{k=1}^{K}\min_{i\leq n}\left(-\sum_{l=1}^{K} x_i^l \lambda^{lk}\right) \leq 0.$$

To do this, we note that for all $\left(\lambda^{kl}\right)_{k,l\leq K}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$, the following are equivalent:

$$\begin{aligned}\left(\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl}\right) + \sum_{k=1}^{K}\min_{i\leq I}\left(-\sum_{l=1}^{K} x_i^l \lambda^{lk}\right) &\leq 0\\ \sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl} &\leq -\sum_{k=1}^{K}\min_{i\leq I}\left(-\sum_{l=1}^{K} x_i^l \lambda^{lk}\right)\\ \sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl} &\leq \sum_{k=1}^{K}\left(\max_{i\leq I}\sum_{l=1}^{K} x_i^l \lambda^{lk}\right).\end{aligned}$$

- **Condition 8:** For every collection of real numbers $\left(\lambda^{kl}\right)_{k,l\leq K}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$, there exists a selection of a state $i\leq I$ per agent $k\leq K$, $i=i_k$, such that

$$\sum_{k=1}^{K}\sum_{l=1}^{K} c^k \lambda^{kl} \leq \sum_{k=1}^{K}\sum_{l=1}^{K} x_{i_k}^l \lambda^{lk}.$$

Proof that Conditions (7) and (8) are equivalent: Given the collection $\left(\lambda^{kl}\right)_{k,l\leq K}$, choose, for $k\leq K$, a state $i=i_k$ such that

$$\sum_{l=1}^{K} x_{i_k}^l \lambda^{lk} = \max_{i\leq I}\sum_{l=1}^{K} x_i^l \lambda^{lk}.$$

Plugging this expression into the last inequality of Condition (7) yields (8), and, conversely, if (8) holds for $i=i_k$ it obviously holds for the maximum as required in (7).

- **Condition 9:** For every $l\leq m$, and collection of real numbers $\left(\lambda^{kl}\right)_{k\leq K}$ such that $\lambda^{kl}\geq 0$ when $k\neq l$, there exists a selection of a state $i=i_l\leq I$, such that

$$\sum_{k=1}^{K} c^k \lambda^{kl} \leq \sum_{r=1}^{K} x_{i_l}^r \lambda^{rl}. \tag{5}$$

Proof that Conditions (8) and (9) are equivalent: If Condition (8) holds, for a given l we can set $\lambda^{kl'} = 0$ for all $l' \neq l$ to obtain Condition (9) for that l (where the index in the right side of (5) was changed from l to r to avoid confusion with the specific l under discussion).

Conversely, assume that Condition (9) holds. Let there be given real numbers $\left(\lambda^{kl}\right)_{k,l\leq K}$ such that $\lambda^{kl} \geq 0$ when $k \neq l$. For each l, consider $\left(\lambda^{kl}\right)_{l\leq K}$ and apply Condition (9) to obtain $i_l \leq I$ such that (5) holds for that l. Condition (8) follows by summation of these inequalities over l to obtain

$$\sum_{l=1}^{K}\sum_{k=1}^{K} c^k \lambda^{kl} \leq \sum_{l=1}^{K}\sum_{r=1}^{K} x_{i_l}^r \lambda^{rl}$$

and then reversing the order of the summation on the left, and on the right replacing the index l by k and the index r by l.

- **Condition 10:**

  (10a) For every collection of non-negative real numbers $\left(\lambda^k\right)_{k\leq K}$, there exists a state $i \leq I$ such that

  $$\sum_{k=1}^{K} c^k \lambda^k \leq \sum_{k=1}^{K} x_i^k \lambda^k.$$

  (10b) For every $l \leq K$, and collection of non-negative real numbers $\left(\lambda^k\right)_{k\leq K, k\neq l}$, there exists a selection of a state $i = i_l \leq I$, such that

  $$\sum_{\substack{k\neq l\\ k\leq K}} \lambda^k \left(c^k - x_{i_l}^k\right) \leq c^l - x_{i_l}^l.$$

Proof that Conditions (9) and (10) are equivalent:

The underlying principle behind the argument is that we can consider Condition (9) for a given l, and separate it into two cases: $\lambda^{ll} \geq 0$ and $\lambda^{ll} < 0$, both allowed by Condition (9). Condition (10a) is equivalent to Condition (9) for every l and every $\left(\lambda^{kl}\right)_k$ with $\lambda^{ll} \geq 0$, whereas Condition (10b) is equivalent to Condition (9) for every l and every $\left(\lambda^{kl}\right)_k$ with $\lambda^{ll} < 0$. The following provides this argument in more detail.

Condition (9) implies Condition (10a):

Let there be given non-negative real numbers $\left(\lambda^k\right)_{k\leq K}$ as in the antecedent of Condition (10a). Pick any l and set $\lambda^{kl} = \lambda^k$ for $k \leq K$. Apply Condition (9) to obtain i_l such that

$$\sum_{k=1}^{K} c^k \lambda^{kl} \leq \sum_{k=1}^{K} x_{i_l}^k \lambda^{kl}$$

and set $i = i_l$ to obtain the consequent of (10a).

Condition (9) implies Condition (10b):

Let there be given $l \leq K$ and non-negative real numbers $\left(\lambda^k\right)_{k\leq K, k\neq l}$ as in the antecedent of Condition (10b). For that l, set $\lambda^{lk} = \lambda^k$ for $k \leq m, k \neq l$, and $\lambda^{ll} = -1$. Apply Condition (9) to $\left(\lambda^{kl}\right)_{k\leq K}$ in order to obtain i_l such that

$$\sum_{k=1}^{K} c^k \lambda^{kl} \leq \sum_{k=1}^{K} x_{i_l}^k \lambda^{kl}.$$

This inequality can be written as (using $\lambda^{ll} = -1$ to move to the second expression)

$$\begin{aligned} \sum_{k=1}^{K} \left(c^k - x_{i_l}^k\right) \lambda^{kl} &\leq 0 \\ \sum_{k\leq K, k\neq l} \left(c^k - x_{i_l}^k\right) \lambda^{kl} &\leq c^l - x_{i_l}^l. \end{aligned}$$

Recalling that $\lambda^{kl} = \lambda^k$ we get the consequent of (10b).

Conditions (10a),(10b) imply Condition (9):

Let there be given $l \leq K$ and real numbers $\left(\lambda^{kl}\right)_{k\leq K}$ such that $\lambda^{kl} \geq 0$ when $k \neq l$, as in the antecedent of Condition (9). Distinguish between two cases:

$\lambda^{ll} \geq 0$. In this case define $\lambda^k = \lambda^{kl}$ (which are all nonnegative) and apply Condition (10a) to obtain the existence of a state i such that

$$\sum_{k=1}^{K} c^k \lambda^k \leq \sum_{k=1}^{K} x_i^k \lambda^k.$$

Denoting $i_l = i$ (and recalling that $\lambda^k = \lambda^{kl}$) we obtain the consequent of Condition (9).

$\lambda^{ll} < 0$. In this case define $\hat{\lambda}^{lk} = \lambda^{kl} / \left|\lambda^{ll}\right|$ so that $\hat{\lambda}^{kl} \geq 0$ for $k \neq l$ and $\hat{\lambda}^{ll} = -1$. Apply Condition (10b) to the state l and the numbers $\left(\lambda^k\right)_{k\leq K, k\neq l}$ defined by $\lambda^k = \hat{\lambda}^{kl} (\geq 0)$. Obtain i_l, such that

$$\sum_{\substack{k\neq l \\ k\leq K}} \lambda^k \left(c^k - x_{i_l}^k\right) \leq c^l - x_{i_l}^l$$

which implies the consequent of Condition (9).

References

[1] James Andreoni and John Miller. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2):737–753, 2002.

[2] Itai Arieli and Manuel Mueller-Frank. Inferring beliefs from actions. *Games and Economic Behavior*, 102:455–461, 2017.

[3] Tilman Börgers. Pure strategy dominance. *Econometrica*, 61(2):423–430, 1993.

[4] Walter Bossert and Kotaro Suzumura. Revealed preference and choice under uncertainty. *SERIEs*, 3:247–258, 2012.

[5] Christopher Chambers and Federico Echenique. *Revealed Preference Theory*. Cambridge University Press, Cambridge, 2016.

[6] Christopher Chambers, Federico Echenique, and Kota Saito. Testing theories of financial decision making. *Proceedings of the National Academy of Science*, 113(15):4003–4008, 2016.

[7] Christopher Chambers, Ce Liu, and Seung-Keun Martinez. A test for risk-averse expected utility. *Journal of Economic Theory*, 163:775–785, 2014.

[8] Federico Echenique and Kota Saito. Savage in the market. *Econometrica*, 83(4):1467–1495, 2015.

[9] F. Edgeworth. *Mathematical Psychics*. Kegan Paul, London, 1881.

[10] Peter C. Fishburn. Separation theorems and expected utility. *Journal of Economic Theory*, 11(1):16–34, 1975.

[11] Itzhak Gilboa. *Theory of Decision under Uncertainty*. Cambridge University Press, Cambridge, 2009.

[12] Itzhak Gilboa, Andrew Postlewaite, Larry Samuelson, and David Schmeidler. Economics: Between prediction and criticism. *International Economic Review*, 59(2):367–390, 2018.

[13] Itzhak Gilboa, M. Rouziou, and O. Sibony. Decision theory made relevant: Between the software and the shrink. *Research in Economics*, 72:240–250, 2018.

[14] Itzhak Gilboa and Larry Samuelson. No-bettingPareto under ambiguity. Technical report, HEC Paris/Tel Aviv Unversity and Yale University, 2020.

[15] Itzhak Gilboa, Larry Samuelson, and David Schmeidler. No-betting Pareto dominance. *Econometrica*, 82(4):1405–1442, 2014.

[16] Hendrik S. Houthakker. Revealed preference and the utility function. *Economica*, 17(66):159–174, 1950.

[17] Felix Kubler, Larry Selden, and Xiao Wei. Asset demand based tests of expected utility maximization. *Journal of Economic Theory*, 104(11):3459–3480, 2014.

[18] Yi-Hsuan Lin. Revealed 'betweenness' preferences over lotteries. Technical report, Intitute of Economics, Academia Sinica, 2019.

[19] Matthew Polisson, John K.-H. Quah, and Ludovic Renou. Revealed preferences over risk and uncertainty. *American Econnomic Review*, 2019. Forthcoming.

[20] Paul A. Samuelson. A note on the pure theory of consumer's behavior. *Economica*, 5(17):61–71, 1938.