

Testable implications of subjective expected utility theory

Eduardo Zambrano

Department of Finance, Mendoza College of Business, University of Notre Dame, South Bend, IN 46556, USA

Abstract

I show that the predictive content of the hypothesis of subjective expected utility maximization critically depends on what the analyst knows about the details of the problem a particular decision maker faces. When the analyst does not know anything about the agent's payoffs or beliefs and can only observe the sequence of actions taken by the decision maker any arbitrary sequence of actions can be implemented as the choice of an agent that solves some intertemporal utility maximization problem under uncertainty.

Keywords: Intertemporal decision theory under uncertainty; Testable implications; Savage–Bayesian rationality

1. Introduction

The purpose of this paper is to answer the following: are there any testable implications of the hypothesis of subjective expected utility maximization (SEU)?

SEU is the theory that states that an agent chooses actions consistent with the maximization of the expectation of a utility function that depends on the action of the agent and

I want to thank David Easley, Larry Blume, Tom Gresik, Dave Kaplan, Todd Keister and Enrique Kawamura for their comments to an earlier version of this work, and especially the associate editor and the referees for making suggestions that substantially improved the quality of the paper. I also received useful comments from the participants at a Cornell seminar. All remaining shortcomings are mine.

on the condition of the environment, and where the expectation is taken over the condition of the environment with respect to some probabilistic belief function. What is shown in this paper is that whether SEU has testable implications crucially depends on what is known by the analyst. In particular, *if (i) the analyst does not know the preferences or the beliefs of the agent and (ii) the analyst can observe the sequence of actions over time chosen by the agent, then SEU has no testable implications.* I show this by providing a SEU representation of a model of intertemporal behavior where the analyst does not know the preferences or the beliefs of the agent. The analyst, however, observes the sequence of actions chosen. In this setup *any* observed behavior can be viewed as the choice of an agent that maximizes expected utility for some utility function, discount factor, uncertainty space and prior belief.

The intuition behind this result is that when choice over time depends on the evolution of a stock that is not known to us then we have enough degrees of freedom in our representation of that stock to interpret *any* observed behavior as the solution to some problem of intertemporal choice. In the case of SEU, the “stock” is the belief held by the agent. If nothing is known about it, it can then be described by the analyst as that which justifies whatever action the agent chose. That such a belief exists and is well-behaved from a probabilistic standpoint arises from the fact that the uncertainty space over which it is defined can also be picked arbitrarily.

The result is not surprising upon reflection about what it says, and to many it is an insight that is known to the research community. Despite this, it is often argued that a departure from SEU is necessary in applications because the behavior it implies seems inconsistent with what agents do in the real world. This suggests the need to make the point clearly by presenting a stark, unambiguous case: one where complete ignorance about preferences and beliefs on the part of the analyst strips SEU of any predictive content. The result of this paper is therefore important because it provides a useful albeit extreme benchmark.

The structure of the rest of the paper is as follows: In Section 2, I present an example of the result. In Section 3, I introduce elements of the theory of statistical decisions. In Section 4, I present the result. Section 5 addresses robustness issues and Section 6 discusses the related literature. Section 7 concludes.

2. A simple example

Assume that the agent has two actions: to carry an umbrella around (a^1) or to leave the umbrella home (a^2). The environment can take one of two conditions: it can either be sunny (S) or rainy (R). The analyst observes the sequence $\{(a_t, y_t)\}_{t=1}^{\infty}$ of actions and conditions of the environment over time, where $(a_t, y_t) \in \{a^1, a^2\} \times \{S, R\} := A \times Y$ and knows nothing about the preferences or the beliefs of the agent. This sequence is the data to be rationalized. Define a 1–1 map between A and Y . This map can be arbitrary, but here I define one that is adapted to the interpretation suggested by the labels of the elements of A and Y : associate a^1 to R and a^2 to S . The utility function that rationalizes this data is based on this map, namely, $u(a^1, R) = u(a^2, S) = 1$ and $u(a^1, S) = u(a^2, R) = 0$. I now turn to the construction of the belief function, which will depend on the evolution of a_t over time and on u . The important thing is that it need not depend at all on the evolution of y_t .

To build this function I adapt an idea used by Oakes (1985). Let $v(y_{t+1} | h_t)$ be the probabilistic belief of the agent about the environment in period $t + 1$ given the history $h_t = \{(a_\tau, y_\tau)\}_{\tau=1}^t$. Define $v(y_{t+1} = S | h_t) = 3/4$ when $a_{t+1} = a^2$ and $v(y_{t+1} = S | h_t) = 1/4$ when $a_{t+1} = a^1$. As a discount factor, pick $\rho = 0$. Notice that, period by period, this agent is choosing a_t , the myopic expected utility maximizing choice when the expectation is taken over y_t with respect to the forecast $v(y_t | h_{t-1})$. The literature on learning in games (cf. Jordan, 1997, p. 154; Nyarko, 1997a, Proposition 7.1) shows that any agent who is best-responding to a prediction rule such as that given by v is, in fact, best responding to a subjective probability distribution on some large parameter space, which in this case we take to be equal to $\{S, R\}^\infty$. Hence the agent is a SEU maximizer (see also the discussion in Section 5).

3. Statistical decision theory

3.1. Actions, the environment

Consider an agent facing an intertemporal choice problem under uncertainty. At each date $t \geq 1$, the agent chooses an action $a_t \in A$. After choosing a_t she observes the condition of the environment $y_t \in Y$. The choice problem is being observed by an analyst who only sees the sequence $\{(a_t, y_t)\}_{t=1}^\infty$. It is therefore important to distinguish what is known by both the agent and analyst from what is known by the agent alone. *The sets A and Y are the primitives of the problem that are given and known to both the agent and the analyst.*

Let A and Y be complete, compact, separable metric spaces with associated σ -fields \mathcal{A} and \mathcal{Y} , respectively. Assume that A has at least two elements and that Y has at least the cardinality of A .

The set of *histories* of length T , \mathcal{H}^T , is the T -fold Cartesian product of $A \times Y$. \mathcal{H}^0 contains the single abstract element h_0 , the null history. The set of all (finite) histories is $\mathcal{H} = \bigcup_{T \geq 0} \mathcal{H}^T$. The set of infinite sequences of profiles (a, y) is denoted \mathcal{Z} . Let the t th coordinate of $z \in \mathcal{Z}$ be z_t and the first t coordinates $z(t)$. Let \mathcal{F} denote the σ -field of subsets of \mathcal{H} derived from the Borel σ -fields on each \mathcal{H}^T . Let $h \cdot h'$ be the concatenation of two histories h and h' . A t -period history will be denoted by h_t . A *strategy* for the agent is a \mathcal{F} -measurable function $\sigma : \mathcal{H} \rightarrow A$ that for each history selects an element of A . Let Σ be the set of strategies of the agent.

The derived elements of the problem that are given and known to both the agent and the analyst are, therefore, A, Y, \mathcal{H}^T ($T = 0, 1, \dots$), $\mathcal{H}, \mathcal{F}, \mathcal{Z}$ and Σ .

3.2. Payoffs, beliefs

Given the agent's choice and the condition of the environment she receives a reward, not observed by the analyst, according to a *payoff function* $u : A \times Y \rightarrow \mathfrak{R}_+$. The agent monitors the condition of the environment according to a collection Θ (of *models of the environment*) and a *prior probability distribution* v defined over Θ . Given one such model θ and a history h the agent has beliefs over the upcoming y given by the map $(h, \theta) \mapsto \eta(\cdot | h, \theta) \in \Delta(Y)$. The interpretation is that the agent considers the condition of

the environment to be a process governed by one of the models in Θ , but is not sure exactly which (hence the prior probability distribution over Θ). Let u be measurable with respect to the product $\mathcal{A} \times \mathcal{Y}$ and assume that Θ is a complete, separable metric space.

The primitives of the problem that are given and known only by the agent are, therefore, u , Θ , v and η .

It will be useful for what follows to derive from η the map $(\sigma, \theta) \mapsto p_\mu(\cdot | \sigma, \theta) \in \Delta(\mathcal{Z})$, a probability distribution over the infinite sequences in \mathcal{Z} .¹ This distribution, derived from the primitives known only to the agent, is consequently only known by the agent.

4. Arbitrary Bayesian rational behavior

A strategy σ is *consistent* with a given sequence $\{(a_t, y_t)\}_{t=1}^\infty$ if $\sigma(h_0) = a_1$ and for all $t \geq 1$, $\sigma(a_1, y_1, \dots, a_t, y_t) = a_{t+1}$. A payoff function $u : A \times Y \rightarrow \mathfrak{R}_+$ is *nontrivial* if there are actions a', a'' and a condition of the environment y such that $u(a', y) \neq u(a'', y)$.

The main result can thus be stated as follows:

Theorem 1. For every sequence $\{(a_t^*, y_t^*)\}_{t=1}^\infty \in \mathcal{Z}$ there are a complete, separable metric space Θ , a prior $v \in \Delta(\Theta)$, a probability function η , a nontrivial utility function u , a discount factor ρ and a strategy σ^* consistent with $\{(a_t^*, y_t^*)\}_{t=1}^\infty$ such that

$$\sigma^* \in \arg \sup_{\sigma \in \Sigma} \int_{\Theta} \int_{\mathcal{Z}} \sum_{t=1}^{\infty} \rho^{t-1} u(z_t) p_\eta(dz | \sigma, \theta) v(d\theta).$$

Proof. Let $f : A \rightarrow Y$ be a 1–1 measurable function.² Define $u(a, y) = 1$ if $y = f(a)$ and $u(a, y) = 0$ otherwise. Then, there is a map $\hat{a} \mapsto \mu_{\hat{a}} \in \Delta(Y)$ such that, for every action $\hat{a} \in A$, $\hat{a} \in \arg \max_{a \in A} \int_Y u(a, y) \mu_{\hat{a}}(dy)$.³

Let $\Theta = Y^\infty$; for every $\theta = (y_1, y_2, \dots)$ and $h_t = (a_1, y_1, \dots, a_t, y_t)$ let $\eta(\cdot | h_t, \theta) = y_{t+1}$. By Kolmogorov's extension theorem the prior v on Θ can be chosen so that $v_1 := \text{marg}_{Y_1} v = \mu_{a_1^*}$ and $v_{t+1} = \text{marg}_{Y_{t+1}} v(\cdot | y_1, \dots, y_t) = \mu_{a_{t+1}^*}$ for every finite sequence (y_1, \dots, y_t) .

Define σ^* as follows: for all $t \geq 0$ and all h_t belonging to a measurable subset of \mathcal{H}^t , $\sigma^*(h_t) = a_{t+1}^*$. Pick any $\rho \in [0, 1)$. Then σ^* maximizes SEU as required.

¹ The derivation of $p_\mu(\cdot | \sigma, \theta)$ is as follows: First, define p_μ recursively for every history $h \in \mathcal{H}$. Let $p_\mu(h_0 | \sigma, \theta) = 1$ and $p_\mu(h \cdot (a, y) | \sigma, \theta) = p_\mu(h | \sigma, \theta) \times \sigma(a | h) \times \mu_\theta(a | h)$. Second, define the *cylinders* $C(h)$ to be the set of paths of play z for which $z(t) = h$, where h belongs to a measurable subset of \mathcal{H}^t . Third, define p_μ over $C(h)$ to be equal to $p_\mu(h | \sigma, \theta)$. This probability measure is then uniquely extended by continuity from the cylinders to the σ -field on \mathcal{Z} generated by the cylinders.

² The existence of this function follows from the fact that A has cardinality of at least two, Y has cardinality at least A and these sets are complete, separable metric spaces.

³ Notice that the distribution $\mu_{\hat{a}}$ can be chosen so that it has full support on Y . For example, $\mu_{\hat{a}} = (1 - \varepsilon)\delta_{f(\hat{a})} + \varepsilon\mu$, where μ has full support and $\varepsilon > 0$ is sufficiently small.

To see this notice that with the structure at hand the problem can be rewritten as

$$\sup_{\sigma \in \Sigma} \int_{\Theta} \sum_{t=1}^{\infty} \rho^{t-1} u(\sigma(h_{t-1}), \theta_t) v(d\theta),$$

which in turn leads to $\sup_{\sigma \in \Sigma} \sum_{t=1}^{\infty} \rho^{t-1} \int_Y u(\sigma(h_{t-1}), y) v_t(dy)$. To show that σ^* solves this problem it suffices (due to the one-shot deviation principle) to check that, given any partial history h_T , there is no expected profitable deviation from $\sigma^*(h_T)$ at date $T + 1$. Recall that $\sigma^*(h_T) = a_{T+1}^*$ and notice that v_t is independent of $\{a_\tau\}_{\tau=1}^T$ for $t \geq T + 1$. Hence, the problem at date $T + 1$ is to choose $a_{T+1} \in A$ to maximize

$$\rho^T \int_Y u(a_{T+1}, y) v_{T+1}(dy) + \sum_{t=T+2}^{\infty} \rho^{t-1} \int_Y u(\sigma^*(h_{t-1}), y) v_t(dy). \quad (1)$$

By construction, the first term in the summation is maximized by $\sigma^*(h_T)$. \square

5. Robustness issues

It is interesting to ponder whether the result holds as we add assumptions about what the analyst knows. For example, suppose the analyst knew the shape of the distribution η , and the parameter space Θ the agent used to represent the uncertainty about the evolution of y_t . Will this invalidate the result? The answer to this question is: not necessarily, as the examples below demonstrate.

- Consider the case presented in Section 2 and assume that the analyst observes $\{(a_t, y_t)\}_{t=1}^{\infty}$, and knows the map η , and the parameter space Θ . The analyst may know the pair (η, Θ) , but nothing is gained if Θ just happened to be $\{S, R\}^{\infty}$ and the agent's beliefs about y_t a convex combination of Dirac measures. Then, as Section 2 shows, any behavior can be rationalized as SEU maximization.
- Consider now the case presented in Section 2 but assume this time that the analyst observes $\{(a_t, y_t)\}_{t=1}^{\infty}$, and that the agent knows that the sequence y_t is *exchangeable* with respect to the agent's prior v . Then, by de Finetti's theorem, with v -probability one, the empirical distribution of y_t converges together with the player's posterior over y_{t+1} to some probability measure $(\mu^*, 1 - \mu^*)$ over $\{R, S\}$. Will this invalidate the result of this paper?

The answer to this question, again, is: not necessarily. For example, consistent with the information given above is the representation with $\Theta = \{\theta\}$, $\mu = (\mu^*, 1 - \mu^*)$ and $v = 1_{\theta}$. This representation, together with the discount factor $\delta = 0$ and the utility function $u(a^1, R) = 1$, $u(a^1, S) = u(a^2, R) = \mu^*$, $u(a^2, R) = 2\mu^*$, rationalizes the given sequence $\{(a_t, y_t)\}_{t=1}^{\infty}$.

- The examples above, when combined, suggest conditions under which the result no longer holds. Consider once more the setup in Section 2 as in the example above when the analyst knew that the agent considered the sequence y_t to be exchangeable. This time the analyst also knows that the agent's representation of the uncertainty over y_t

is given by μ over $\{R, S\}$ with $\Theta = [0, 1]$ for some non-singular prior ν over Θ . Exchangeability implies that posterior beliefs are, in the limit, constrained by the realization of θ , which means that the utility function that rationalizes a particular sequence $\{(a_t, y_t)\}_{t=1}^{\infty}$, cannot be independent of θ . As a consequence, there is no state independent utility function that rationalizes $\{(a_t, y_t)\}_{t=1}^{\infty}$ in this case.

6. Related literature

In one of the first attempts to understand the restrictions imposed by SEU, Pearce (1984) showed that when the analyst knows the agent's preferences over mixtures of actions SEU rules out players choosing strictly dominated strategies. Ledyard (1986) and Börgers (1993) assume that only preferences over pure strategy outcomes are observable, and show that, as in Pearce (1984), a notion of domination carries the testable content of SEU. Blume and Easley (1998) and Nyarko (1997b) have shown that, given an observable profile of utility functions any stochastic process of undominated actions can be the outcome of a model of intertemporal optimization and Bayesian learning. Nyarko (1997b) considers the zero discount factor case while Blume and Easley (1998) consider the more general case. *From the methodological standpoint*, Blume and Easley (1998) *can be thought of as a direct precursor of the present work*.

Lo (2000) shows that if one is restricted to the observation of only one choice from a finite set of acts, the subjective expected utility model is observationally indistinguishable from all models of preference that satisfy Savage's axiom P3, which is a form of monotonicity. Epstein (2000) points to the need for the analyst to observe choices from at least two different sets of choices where the agent has the same belief if one is to be able to refute SEU. Border (1992) assumes that the analyst is able to observe the entire choice function of the agent and shows that any choice function consistent with SEU must not be stochastically dominated.

Another closely related paper is that of Green and Park (1996), which asks whether a strategy can be rationalized by maximization of conditional state dependent utility. They identify a necessary and sufficient condition, in an environment with a correctly specified prior, for a strategy to be rationalizable. Their condition requires for a plan not to react to "irrelevant" information. It is a very weak condition.

7. Conclusions

The present paper contributes to the literature associated with how little restrictions the assumptions of rationality impose on individual and collective behavior by showing that when the analyst knows nothing about the preferences or beliefs of an agent any sequence of actions observed by the analyst can be the outcome of some model of intertemporal optimization and Bayesian learning.

Not all predictive content of SEU is lost in practice, however, because auxiliary assumptions about what is known to the analyst can be made which, jointly with observability of actions and SEU, generate testable implications. In this sense, the result presented in this

paper is an instance of a principle known to modern philosophers of science: that no core set of theoretical assumptions can be contradicted in isolation (Caldwell, 1982, Chapter 4).

References

- Blume, L., Easley, D., 1998. Rational expectations and rational learning. In: Majumdar, M. (Ed.), *Organizations with Incomplete Information*. Cambridge Univ. Press., Cambridge, pp. 61–109.
- Border, K., 1992. Revealed preference, stochastic dominance, and the expected utility hypothesis. *J. Econ. Theory* 56, 20–42.
- Börgers, T., 1993. Pure strategy dominance. *Econometrica* 61, 423–430.
- Caldwell, B., 1982. *Beyond Positivism: Economic Methodology in the Twentieth Century*. Allen & Unwin, London.
- Epstein, L., 2000. Are probabilities used in markets? *J. Econ. Theory* 91, 86–90.
- Green, E., Park, I., 1996. Bayes contingent plans. *J. Econ. Behav. Organ.* 31, 225–236.
- Jordan, J., 1997. Bayesian learning in games: a non-Bayesian perspective. In: Bicchieri, C., Jeffrey, R., Skyrms, B. (Eds.), *The Dynamics of Norms*. Cambridge Univ. Press, Cambridge, pp. 149–174.
- Ledyard, J., 1986. The scope of the hypothesis of Bayesian equilibrium. *J. Econ. Theory* 39, 59–82.
- Lo, K., 2000. Rationalizability and the Savage axioms. *Econ. Theory* 15, 727–733.
- Nyarko, Y., 1997a. Savage–Bayesian agents play a repeated game. In: Bicchieri, C., Jeffrey, R., Skyrms, B. (Eds.), *The Dynamics of Norms*. Cambridge Univ. Press, Cambridge, pp. 175–197.
- Nyarko, Y., 1997b. Convergence in economic models with Bayesian hierarchies of beliefs. *J. Econ. Theory* 74, 266–296.
- Oakes, D., 1985. Self-calibrating priors do not exist. *J. Amer. Statist. Assoc.* 80, 339.
- Pearce, D., 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1050.