

On the Rejectability of the Subjective Expected Utility Theory*

Konrad Grabiszewski[†]

2016

Abstract

State space, a key element of the Subjective Expected Utility (SEU) theory, is not observable. This implies that, in order to test the SEU theory, it is necessary to assume some state space. Consequently, if the SEU theory is rejected, then it is appropriate to conduct a robustness check; that is, to search for a different state space and a probability over that state space which together do not lead to the rejection of the SEU theory. To find such state space and probability means to SEU-rationalize the agent's behavior. I show how to conduct the process of SEU-rationalization and determine when an SEU-rationalization is possible.

Keywords: SEU theory; state space; SEU-rationalization

JEL classification: D01; D03; D81.

*For their valuable comments, I would like to thank Pritha Dev, Umberto Garfagnini, Itzhak Gilboa, Andrei Gomberg, Alex Horenstein, Jeffrey Kline, Jean-Philippe Lefort, Erkut Ozbay, Levent Ülkü, Burkhard Schipper, Tridib Sharma, Joel Sobel, Peter Wakker, the anonymous reviewer, and the participants of the FUR XV International Conference, 14th SAET Conference, workshops organized by HEC (Paris), and ITAM, and seminar at University of Miami.

[†]Address: Department of Economics, University of Miami, Coral Gables, FL 33124, USA; konrad.grabiszewski@gmail.com

1 Introduction

Like the majority of theories in the literature on uncertainty/ambiguity, the Subjective Expected Utility (SEU) theory developed by Leonard J. Savage (Savage (1972)) begins with a state space Ω with a generic element ω , and a set of consequences X with a generic element x . Then, there is a set of acts (functions from Ω to X) \mathcal{F} with a generic element f . A binary order \succsim is defined over \mathcal{F} , and the axioms imposed on \succsim guarantee the existence of an SEU representation of \succsim . To test the SEU theory means to verify whether \succsim satisfies the Savage axioms.

Unfortunately, as it has been noted in the literature,¹ the state space is not observable and, consequently, is not part of the researcher's (hereafter, referred to as "he") data. This implies that the researcher is unable to test the SEU theory because with no state space there is no set of acts, and with no set of acts there is no binary order to be tested.

If the researcher assumes a specific choice theory, then he might be able to derive the state space from the agent's (hereafter, referred to as "Ann" or "she") observable choices. For example, Proposition 3 in Schipper (2013) might help to conduct such an exercise under the assumption that the agent satisfies the SEU theory. But assuming a theory to derive the agent's state space, and then using that state space to test the very theory that has just been assumed is, obviously, an inappropriate approach.

In what I call the *standard approach* to testing the SEU theory, the researcher implicitly *assumes* some state space Ω . But what state space should he assume? Economic theory provides no answer to this question, a problem Savage already noticed sixty years ago: "I am unable to formulate criteria for selecting these small worlds and indeed believe that their selection may be a matter of judgment and experience about which it is impossible to enunciate complete and sharply defined general principles" (Savage (1972, p. 16)).

If the researcher follows the standard approach and rejects the SEU theory, then his conclusion is not unconditionally valid. Rather, his conclusion depends on the state space Ω that he assumed at the start, and it would be more appropriate to say that the SEU theory is rejected *conditional* on the state space being Ω . It is possible that there exist a different state space $\tilde{\Omega}$ and a probability measure $\tilde{\lambda}$ over $\tilde{\Omega}$ which together do *not* lead to the rejection of the SEU theory.

For example, consider the two-urn Ellsberg experiment. With a typically assumed state

¹See Billot and Vergopoulos (2014), Blume et al. (2006), Blume et al. (2009), Epstein (2010), Gilboa (2009), Gilboa and Schmeidler (2004), Grant et al. (2015), Karni (2008), and Machina (2003).

space, this experiment rejects the SEU theory. However, Kadane (1992) shows that the experiment fits within the SEU framework as long as the agent assigns an arbitrarily small probability to the experimenter being malevolent. It does not seem unreasonable to assume that, from the agent’s perspective, the experimenter is malevolent. After all, there is nothing in this experiment’s design that prevents the agent from having such a belief. Hence, with a different state space, the two-urn Ellsberg experiment does not reject the SEU theory.

Since the state space Ω is assumed rather than derived (from data or theory), it is only appropriate to conduct a robustness check. That is, the researcher should verify whether it is possible to construct a new state space $\tilde{\Omega}$ and a probability $\tilde{\lambda}$ over $\tilde{\Omega}$ in such a way that the agent’s choices appear to be in accordance with the SEU theory. If the researcher constructs such $\tilde{\Omega}$ and $\tilde{\lambda}$, then I say that he is able to *SEU-rationalize* the agent’s behavior. The standard approach does not include any robustness check. In this paper, I fill the gap by showing how to conduct the process of SEU-rationalization and determining when an SEU-rationalization is possible.

In section 2, I present and analyze the process of SEU-rationalization. In section 3, I apply my procedure to the classical three-color Ellsberg experiment, which is the most popular example of a violation of the SEU theory. I show how to SEU-rationalize the behavior observed in this experiment. Section 4 discusses the relevant literature. Section 5 concludes.

2 SEU-rationalization

I consider the researcher who follows the standard approach to testing the SEU theory, rejects the SEU theory, and would like to conduct a robustness check. The quadruple $(\Omega, X, \mathcal{F}, \succsim)$ —a state space Ω , a set of consequences X , a set of acts \mathcal{F} , and a binary order \succsim over \mathcal{F} —is taken as given and constitutes the researcher’s benchmark data set.

The objective of SEU-rationalization is to translate the original $(\Omega, X, \mathcal{F}, \succsim)$ into a new $(\tilde{\Omega}, \tilde{X}, \tilde{\mathcal{F}}, \tilde{\succsim})$ and construct a probability measure $\tilde{\lambda}$ over $\tilde{\Omega}$ such that $\tilde{\succsim}$ has an SEU representation with respect to $\tilde{\lambda}$. This new $(\tilde{\Omega}, \tilde{X}, \tilde{\mathcal{F}}, \tilde{\succsim})$ is a re-interpretation of the original $(\Omega, X, \mathcal{F}, \succsim)$. Obviously, allowing for any $(\tilde{\Omega}, \tilde{X}, \tilde{\mathcal{F}}, \tilde{\succsim})$ turns the procedure of SEU-rationalization into a trivial and uninteresting exercise. I allow to change only the state space. This implies that several conditions must be satisfied.

First, I require to keep the same set of consequences (condition 1 in Definition 2.1). Second, \mathcal{F} and \succsim must be consistently translated into $\tilde{\mathcal{F}}$ and $\tilde{\succsim}$, respectively (conditions 2 and 3 in

Definition 2.1).

Consistency between $\tilde{\mathcal{F}}$ and \mathcal{F} means that the perception of acts does not change. Since there is a new state space $\tilde{\Omega}$, there is a new set of acts $\tilde{\mathcal{F}}$ with a generic element $\tilde{f} : \tilde{\Omega} \rightarrow X$. Take f that is translated into \tilde{f} . I require that \tilde{f} and f are perceived in the same way. Observe that f and \tilde{f} are defined on different domains. Hence, in order to compare f and \tilde{f} it is necessary to look at their common codomain X . Let X_f , a subset of the range of f , be the collection of all possible consequences associated with act f . If $x \in X_f$, then the agent perceives that such a consequence can be obtained if she chooses f . I assume that the researcher's data set includes one X_f for each $f \in \mathcal{F}$. Note that \tilde{f} and $\tilde{\lambda}$ generate the probability measure $\tilde{\lambda}_f$ on X . The support of that measure is the collection of all possible consequences associated with act \tilde{f} . I say that f is consistently translated into \tilde{f} if the support of $\tilde{\lambda}_f$ is precisely X_f .

Condition 3 in Definition 2.1 requires that $\tilde{\succsim}$ be consistent with \succsim ; that is, $\tilde{f} \tilde{\succsim} \tilde{g}$ if and only if $f \succsim g$. This implies that the ranking of acts does not change.

Let \succ and \sim denote the strict and indifference part of \succsim , respectively. Let f_x be a constant act that generates x (i.e., $f(\omega) = x$ for each ω). Let $f_{\max f}$ denote the constant act that yields the best possible consequence associated with f . Formally, $\max f \in X_f$ and $f_{\max f} \succsim f_x$ for each $x \in X_f$. Let $f_{\min f}$ be defined in a similar way.

Definition 2.1 *SEU-rationalization.*

Fix a quadruple $(\Omega, X, \mathcal{F}, \succsim)$. For each $f \in \mathcal{F}$, fix a set X_f . To SEU-rationalize a binary order \succsim means to construct a quadruple $(\tilde{\Omega}, \tilde{X}, \tilde{\mathcal{F}}, \tilde{\succsim})$ and a probability measure $\tilde{\lambda}$ over $\tilde{\Omega}$ such that:

1. $\tilde{X} = X$.
2. $\tilde{\lambda}$ and \tilde{f} generate a probability measure on X such that its support is X_f .
3. $\tilde{f} \tilde{\succsim} \tilde{g}$ if and only if $f \succsim g$.
4. $\tilde{\succsim}$ has an SEU representation with respect to $\tilde{\lambda}$. That is, there exists $u : X \rightarrow \mathbb{R}$ such that $\tilde{\succsim}$ is represented by $\tilde{V} : \tilde{\mathcal{F}} \rightarrow \mathbb{R}$, defined as $\tilde{V}(\tilde{f}) := \sum_{\tilde{\omega} \in \tilde{\Omega}} u(\tilde{f}(\tilde{\omega})) \tilde{\lambda}(\tilde{\omega})$.

Below, I present the assumptions imposed on the triple $(X, \mathcal{F}, \succsim)$ which allow to SEU-rationalize a given binary order \succsim . There are no restrictions on Ω because it is the state space that is to be changed.

Assumption 1

1. X and \mathcal{F} are finite.
2. \succsim is a complete and transitive binary order over \mathcal{F} .
3. For each $x \in X$, there is $f_x \in \mathcal{F}$.

Assumption 2

For each $f \in \mathcal{F}$, if $f_{\max f} \succ f_{\min f}$, then $f_{\max f} \succ f \succ f_{\min f}$.

Assumption 1.1 is a rather natural assumption given that data sets are finite. Nevertheless, in the appendix, I replicate my result allowing for infinite sets of consequences and acts. Assumption 1.2, combined with Assumption 1.1, guarantees the existence of utility representation of \succsim , V . The purpose of having constant acts (Assumption 1.3) is, obviously, to construct a utility function u on the set of consequences; $u(x) := V(f_x)$. Note that Assumption 2 can be seen as an extension of the Simple Dominance Axiom introduced by Barberà (1977), which holds that $\{x\} \succ \{y\}$ implies $\{x\} \succ \{x, y\} \succ \{y\}$.

The following proposition is the main result of this paper.

Proposition 1

Fix a quadruple $(\Omega, X, \mathcal{F}, \succsim)$. For each $f \in \mathcal{F}$, fix a set X_f . If Assumptions 1 and 2 hold, then an SEU-rationalization of \succsim is possible.

I prove Proposition 1 by construction. I develop a two-stage procedure of SEU-rationalization and show that a given binary order \succsim can be SEU-rationalized if Assumptions 1 and 2 are satisfied.

First stage (from acts to options). In the first stage, transform the set of acts \mathcal{F} into the set of options, denoted by $\hat{\mathcal{F}}$. An option \hat{f} is defined as a pair (f, X_f) , where f is the act (the option's name) and X_f is the collection of consequences associated with act f .

To understand why f is transformed into (f, X_f) rather than X_f alone, consider a bet on the next winner of the U.S. presidential elections. There are two acts: a Democrat, denoted by f_D , and a Republican, denoted by f_R . Suppose that picking the winner yields \$100, while picking the loser costs the agent \$50. Each act is associated with the same set of consequences, that is, $X_{f_D} = X_{f_R} = \{\$100, -\$50\}$. If the acts were to be represented only

by the associated sets of consequences, then f_D and f_R would be perceived as the same object, or it would be necessary to assume that the agent is indifferent between two acts whenever their associated sets of consequences are the same. As the betting example suggests, this is generally not a correct approach; hence, a name is included in the definition of an option.

Next, construct the binary order $\hat{\succsim}$ on $\hat{\mathcal{F}}$ in a consistent way from \succsim : $\hat{f} \hat{\succsim} \hat{g}$ if and only if $f \succsim g$. Let $\hat{\succ}$ and $\hat{\sim}$ denote the strict and indifference part of $\hat{\succsim}$, respectively. Let \hat{f}_x be an option generated by a constant act f_x .

Note that constructing $(X, \hat{\mathcal{F}}, \hat{\succsim})$ from $(X, \mathcal{F}, \succsim)$ is always possible; that is, the assumptions presented above are irrelevant to accomplish the first stage of SEU-rationalization. However, they are important in the second stage.

Second stage (from options to acts). The second stage begins with the triple $(X, \hat{\mathcal{F}}, \hat{\succsim})$ as derived in the first stage. First, construct a utility function \hat{V} defined over $\hat{\mathcal{F}}$, which exists because of Assumption 1. This assumption also makes it possible to uniquely define a utility function u over X as $u(x) := \hat{V}(\hat{f}_x)$.

For each option \hat{f} , construct a probability measure $\hat{\lambda}_f$ defined over X such that (i) the support of $\hat{\lambda}_f$ is X_f , and (ii) \hat{V} can be reconstructed as $\hat{V}(\hat{f}) = \sum_{x \in X_f} u(x) \hat{\lambda}_f(x)$. The construction of $\hat{\lambda}_f$ s is based on the repetitive use of the Intermediate Value Theorem. Take \hat{f} such that $X_f = \{x_1, \dots, x_K\}$. Without loss of generality, assume that $\hat{f}_{x_1} \hat{\succ} \dots \hat{\succ} \hat{f}_{x_K}$. Assumption 2 indicates that $\hat{f}_{x_1} \hat{\succ} \hat{f} \hat{\succ} \hat{f}_{x_K}$. However, because of Assumption 1, it is known how the agent ranks \hat{f} in comparison with each \hat{f}_{x_n} for $n = 1, \dots, K$. Next, without loss of generality, assume that $\hat{f}_{x_1} \hat{\succ} \hat{f} \hat{\succ} \hat{f}_{x_2}$. Hence, it is true that $\hat{f}_{x_1} \hat{\succ} \hat{f} \hat{\succ} \hat{f}_{x_n}$ for each $n = 2, \dots, K$. The Intermediate Value Theorem says that, for each $n = 2, \dots, K$, there is an $\alpha_n \in (0, 1)$ such that $\hat{V}(\hat{f}) = u(x_1)\alpha_n + u(x_n)(1 - \alpha_n)$. Now, take $\beta_2, \dots, \beta_K \in (0, 1)$ such that $\beta_2 + \dots + \beta_K = 1$, and observe that $\beta_2[u(x_1)\alpha_2 + u(x_2)(1 - \alpha_2)] + \dots + \beta_K[u(x_1)\alpha_K + u(x_K)(1 - \alpha_K)] = \hat{V}(\hat{f})$. This gives $\hat{V}(\hat{f}) = \sum_{x \in X_f} u(x) \hat{\lambda}_f(x)$, where $\hat{\lambda}_f(x_1) = \beta_2\alpha_2 + \dots + \beta_K\alpha_K$ and, for $n = 2, \dots, K$, $\hat{\lambda}_f(x_n) = \beta_n(1 - \alpha_n)$. Each element of the set $\hat{\lambda}_f = \{\hat{\lambda}_f(x_1), \dots, \hat{\lambda}_f(x_K)\}$ is a number in $(0, 1)$, and their sum is one. That is, $\hat{\lambda}_f$ denotes a desired probability.

The state space $\tilde{\Omega}$ is defined as a product of X_f s. That is, $\tilde{\Omega} = \times_{f \in \hat{\mathcal{F}}} X_f$. Defining a state space in this way implies that a state is a function from the set of acts to the set of consequences (see Gilboa (2009), Karni and Schmeidler (1991), Karni and Vierø (2013), and Schmeidler and Wakker (1987)). The probability measure $\tilde{\lambda}$ over $\tilde{\Omega}$ is built as a product measure of the probability measures $\hat{\lambda}_f$ s. The set of new acts $\tilde{\mathcal{F}}$ is derived from the set of options $\hat{\mathcal{F}}$ by defining $\tilde{f} : \tilde{\Omega} \rightarrow X$ as a projection function. The binary order $\tilde{\succsim}$ on $\tilde{\mathcal{F}}$ is constructed in a consistent way from $\hat{\succsim}$: $\tilde{f} \tilde{\succsim} \tilde{g}$ if and only if $\hat{f} \hat{\succsim} \hat{g}$. Finally, define $\tilde{V} : \tilde{\mathcal{F}} \rightarrow \mathbb{R}$

as $\tilde{V}(\tilde{f}) := \sum_{\tilde{\omega} \in \tilde{\Omega}} u(\tilde{f}(\tilde{\omega}))\tilde{\lambda}(\tilde{\omega})$.

Since $\tilde{\lambda}$ is a product measure and \tilde{f} is a projection function, they induce a probability measure on X that is the same as $\hat{\lambda}_f$ (condition 2 in Definition 2.1). Note that $\tilde{\succsim}$ is constructed in a consistent way from \succsim (condition 3 in Definition 2.1). By construction of $\tilde{\lambda}$ and \tilde{f} , it follows that $\tilde{V}(\tilde{f}) = \hat{V}(\hat{f})$. This implies that $\tilde{\succsim}$ has an SEU representation (condition 4 in Definition 2.1). This ends the proof of Proposition 1.

2.1 Analysis of Assumption 2

It is crucial that Assumption 2 be formulated in terms of \succ , instead of \succsim . Take f with $X_f = \{x, y\}$ where $f_x \succ f_y$. If $f \sim f_y$, then the agent must assign probability 1 to y . However, this would violate the fact that the agent considers both x and y to be possible when she chooses f .

In order to better understand the limits imposed by Assumption 2, consider the following two possible violations. First, suppose that the agent prefers $f_{\min f}$ to f . She does not want to try to improve her well-being, even though there are no costs involved in attempting this. After all, if she chooses f , the worst possible result is $\min f$. Next, consider the agent who prefers f to $f_{\max f}$. Instead of choosing the best possible consequence that f can generate, which is $\max f$, the agent picks f .

Although Assumption 2 may appear to be an obvious and harmless requirement, I argue that reasonable violations of Assumption 2 are possible. Consider the agent's problem concerning her choice of destination for her next summer vacation. The consequences represent different levels of satisfaction derived from a summer trip: x_1 (good) and x_2 (bad). The acts are destinations: g (Georgia), h (Hawaii), and i (Idaho), with $X_g = \{x_1\}$, $X_h = \{x_1, x_2\}$, and $X_i = \{x_2\}$. I consider two agents, Ann and Bob. Deciding between Hawaii and Georgia, Ann chooses to go to Hawaii which she has never visited. Since Ann strictly prefers Hawaii over Georgia, her behavior violates Assumption 2. This choice may appear to be puzzling. However, visiting Hawaii means not only having a vacation but also learning what a vacation in Hawaii means. Ann evaluates her acts by their weighted sum, $V(f) = EU(f) + \beta|X_f|$, where $EU(f)$ is the expected utility derived from f and $|X_f|$ is a cardinality of X_f . (In this context, cardinality can be interpreted as a measure of how well the agent knows f .) A parameter β measures the importance of learning. Ann likes to learn, and her β is strictly bigger than zero. Alternatively, consider Bob, who prefers Idaho over Hawaii. For Bob, going to Hawaii will guarantee an expected value that is at least the same as a trip to Idaho.

However, Bob's β is negative because he dislikes learning and prefers to stick to what he already knows best. In consequence, he, like Ann, violates Assumption 2.

3 SEU-rationalization of the three-color Ellsberg experiment

Consider the three-color Ellsberg experiment. In the urn, there are three balls, and only one of them is red. Each of the remaining two can be either green or blue. However, the number of green balls in the urn is unknown. One ball will be drawn from the urn. The typically assumed state space is $\Omega = \{R, G, B\}$, where the letter indicates which ball will be drawn, and the set of consequences X consists of two elements, S (uccess) and F (ailure). The agent has two decisions to make: choose between a bet on a red ball f_r and a bet on a green ball f_g (the agent picks f_r), and choose between a bet on a red or blue ball f_{rb} and a bet on a green or blue ball f_{gb} (the agent picks f_{gb}).

$$f_r \succ f_g \quad \text{and} \quad f_{gb} \succ f_{rb} \tag{1}$$

Assuming that the state space is $\Omega = \{R, G, B\}$ implies that the binary order depicted in (1) is inconsistent with the SEU theory (Table 1).

	f_r	f_g	f_{gb}	f_{rb}	probability
R	S	F	F	S	—
G	F	S	S	F	—
B	F	F	S	S	—

Table 1

The procedure of SEU-rationalization requires, first, to take Table 1 and translate it into the set of options. This yields \hat{f}_r , \hat{f}_g , \hat{f}_{gb} , and \hat{f}_{rb} such that $X_r = X_g = X_{rb} = X_{gb} = \{S, F\}$. Define \hat{V} and u in the way that agrees with (1): $\hat{V}(\hat{f}_r) = 3$, $\hat{V}(\hat{f}_g) = 1$, $\hat{V}(\hat{f}_{gb}) = 4$, $\hat{V}(\hat{f}_{rb}) = 2$, $u(S) = 5$, and $u(F) = 0$. Define the probabilities associated with each option: $\hat{\lambda}_{f_r}(S) = 0.6$, $\hat{\lambda}_{f_g}(S) = 0.2$, $\hat{\lambda}_{f_{gb}}(S) = 0.8$, and $\hat{\lambda}_{f_{rb}}(S) = 0.4$.

Next, construct the state space $\tilde{\Omega} = \{S, F\} \times \{S, F\} \times \{S, F\} \times \{S, F\}$. Each state consists of four coordinates. Also, construct the probability $\tilde{\lambda}$ over $\tilde{\Omega}$ as the product measure of $\hat{\lambda}_{f_r}$,

$\hat{\lambda}_{f_g}$, $\hat{\lambda}_{f_{gb}}$, and $\hat{\lambda}_{f_{rb}}$. Finally, transform the set of options into the set of new acts by defining the new acts as projection functions. For example, option \hat{f}_r becomes act $\tilde{f}_r : \tilde{\Omega} \rightarrow X$, defined as $\tilde{f}_r(\tilde{\omega}) = S$ if and only if the first coordinate of $\tilde{\omega}$ is S . This yields Table 2.

	\tilde{f}_r	\tilde{f}_g	\tilde{f}_{gb}	\tilde{f}_{rb}	probability		\tilde{f}_r	\tilde{f}_g	\tilde{f}_{gb}	\tilde{f}_{rb}	probability
$\tilde{\omega}_1$	S	S	S	S	0.0384	$\tilde{\omega}_9$	F	S	S	S	0.0256
$\tilde{\omega}_2$	S	S	S	F	0.0576	$\tilde{\omega}_{10}$	F	S	S	F	0.0384
$\tilde{\omega}_3$	S	S	F	S	0.0096	$\tilde{\omega}_{11}$	F	S	F	S	0.0064
$\tilde{\omega}_4$	S	S	F	F	0.0144	$\tilde{\omega}_{12}$	F	S	F	F	0.0096
$\tilde{\omega}_5$	S	F	S	S	0.1536	$\tilde{\omega}_{13}$	F	F	S	S	0.1024
$\tilde{\omega}_6$	S	F	S	F	0.2304	$\tilde{\omega}_{14}$	F	F	S	F	0.1536
$\tilde{\omega}_7$	S	F	F	S	0.0384	$\tilde{\omega}_{15}$	F	F	F	S	0.0256
$\tilde{\omega}_8$	S	F	F	F	0.0576	$\tilde{\omega}_{16}$	F	F	F	F	0.0384

Table 2

Table 2 has sixteen rows, with each denoting a state. Columns \tilde{f}_r , \tilde{f}_g , \tilde{f}_{gb} , and \tilde{f}_{rb} show the realizations of bets at a given state. For example, at $\tilde{\omega}_1$ all bets yield S . The probability column depicts the probability that a given state will occur.

However, Table 2 allows for states which are not reasonable. For example, it is not reasonable that a ball is simultaneously red and green (state $\tilde{\omega}_1$). In the suggested SEU-rationalization, there are 13 unreasonable states and 3 reasonable states. The three reasonable states are $\tilde{\omega}_7$ (i.e., red ball is drawn), $\tilde{\omega}_{10}$ (i.e., green ball is drawn), and $\tilde{\omega}_{13}$ (i.e., blue ball is drawn). These are the states R , G , and B , respectively, from the state space Ω (see Table 1). The total mass assigned to these three states is only 0.1792. For this reason, this SEU-rationalization is not satisfactory. Therefore, the obvious question to ask is: What is the smallest probability assigned to the unreasonable states that would still allow for the SEU-rationalization of the Ellsberg experiment? The answer is: As small as desired. Let $\varepsilon < 0.3$ be the total mass assigned to unreasonable states.

	\tilde{f}_r	\tilde{f}_g	\tilde{f}_{gb}	\tilde{f}_{rb}	probability
R	S	F	F	S	$\frac{1}{10} - \frac{1}{3}\varepsilon$
G	F	S	S	F	$\frac{1}{10} - \frac{1}{3}\varepsilon$
B	F	F	S	S	$\frac{8}{10} - \frac{1}{3}\varepsilon$
Z	S	F	S	F	$\varepsilon > 0$

Table 3

In Table 3, there is only one unreasonable state Z . It is easy to verify that this SEU-rationalization captures the choices in (1): $\tilde{V}(\tilde{f}_r) = \frac{5}{10} + \frac{10}{3}\varepsilon > \tilde{V}(\tilde{f}_g) = \frac{5}{10} - \frac{5}{3}\varepsilon$; and $\tilde{V}(\tilde{f}_{gb}) = \frac{45}{10} + \frac{5}{3}\varepsilon > \tilde{V}(\tilde{f}_{rb}) = \frac{45}{10} - \frac{10}{3}\varepsilon$.

The existence of state Z in Table 3 could be due to the agent's being unaccustomed to decision problems expressed in terms of an urn and some balls. As Gilboa (2009) writes, "David Schmeidler often says, 'Real life is not about balls and urns.' Indeed, important decisions involve war and peace, recessions and booms, diseases and cures." An ε -small probability assigned to Z can be interpreted as a manifestation of the agent's misunderstanding the problem or miscalculating the probability.

If the experiment's subjects are college students, which is often the case in economics experiments, then the analysis conducted by psychologists suggests that such a misunderstanding or miscalculation is not unlikely. For example, Standing (2006) finds that only 33 percent of undergraduate students at one liberal arts college achieved a perfect score in a test that focused on very elementary skills: the hardest task was to compute 92×32 (and students were permitted to use unlimited time and write a draft work). In Standing et al. (2006), a similar test was applied to undergraduate students majoring in business and economics. The hardest question was to compute $(36 \times 7) + (33 \times 7)$, and only 40 percent of subjects answered all questions correctly. From LeFevre et al. (2014), we learn that basic arithmetic skills among college students have been declining from 1993 to 2005. If the students struggle with adding or multiplying natural numbers, it only seems natural to expect that they make mistakes in computing subjective probabilities.

A typical explanation of the three-color Ellsberg experiment allows only for the reasonable states (i.e., $\Omega = \{R, G, B\}$) but assumes that the agent behaves as if following a mathematically sophisticated theory (at least, more sophisticated than the SEU theory). Here, I suggest an alternative explanation: the agent behaves as if following the SEU theory but she

assigns an arbitrarily small probability to an unreasonable state (Table 3). Which explanation of the agent’s behavior is “better”? Unfortunately, there is no clear answer. Gilboa and Samuelson (2012) observe that in situations like this, “people typically bring subjective criteria to bear in making this choice (i.e., choosing among theories), tending to select theories that seem a priori reasonable, intuitive, simple, elegant, familiar, or that satisfy a variety of other considerations.” Consequently, in the case of the three-color Ellsberg experiment, the reader is left to use subjective criteria to decide which explanation seems “better.”

However, I need to stress that from the empirical perspective—that is, from a perspective that bases conclusions only on observable data—the Ellsberg experiment does not reject the SEU theory. Although it is possible to test whether or not Z is a Savage-null event (and if it is, then the agent does not assign even an ε -small probability to Z), such a test does not contradict my claim. Note that when I use the phrase *the Ellsberg experiment*, I am referring to an experiment consisting of the four alternatives and two choices described above in (1). For clarity, I call this *the original Ellsberg experiment*. Testing whether or not Z is Savage-null requires that the researcher expand the experiment by adding new alternatives and observing more than two choices. This would not be *the original Ellsberg experiment* (which my claim is about) but rather *the extended Ellsberg experiment*.

Suppose that the extended Ellsberg experiment shows that not only Z but also all the remaining unreasonable states are Savage-null. This would indicate that the procedure of SEU-rationalization proposed in this paper cannot defend the SEU theory against the evidence provided by the original Ellsberg experiment. This also indicates how to design a robust experiment to test the SEU theory: in order to survive the procedure of SEU-rationalization, an experiment that is to reject the SEU theory must be nested in a bigger experiment which rules out all states that can be used by SEU-rationalization.

4 Related literature

What differentiates my paper from the literature on choice theory under uncertainty/ambiguity is that my ambition is not to introduce a new theory, but rather to develop a procedure which allows the researcher to not reject the SEU theory. One of the main objectives of the theories is to have the unique representation of the agent’s belief (additive probability, non-additive probability, set of probabilities). In contrast, neither the state space $\tilde{\Omega}$ nor the probability $\tilde{\lambda}$ constructed to SEU-rationalize the agent’s behavior is unique. There could be another state space Ω' and another probability λ' that would also do the job. In fact, even for a

given $\tilde{\Omega}$, the probability $\tilde{\lambda}$ is not unique because there could be a distinct $\tilde{\lambda}^*$ such that $\tilde{\Omega}$ and $\tilde{\lambda}^*$ also lead to an SEU-rationalization. The non-uniqueness is a result of constructing *both* a state space and a probability rather than *only* a probability. However, the lack of uniqueness is not a concern in this paper. This is because the main goal is to find *some* state space with *some* probability such that the agent appears to be following the SEU theory. In order to emphasize that I am not constructing a new theory, I have been employing the term *assumption*, rather than the term *axiom*.

Two papers which are closely related to this paper are Gilboa and Schmeidler (1994) and Lipman (1999). Gilboa and Schmeidler (1994) considers an agent represented by the Choquet Expected Utility (Schmeidler (1989)) and shows that “the non-additivity of the “probability” v may be explained by “omitted” states of the world. If those were introduced into the model explicitly, the non-additivity would disappear.”. The first difference between my paper and Gilboa and Schmeidler (1994) is that my key requirement (Assumption 2) is weaker than assuming the Choquet Expected Utility theory. In addition, the state space $\tilde{\Omega}$ in my paper is constructed from the subsets of the set of consequences, while in Gilboa and Schmeidler (1994) the state space is an extension of a pre-assumed state space. Lipman (1999) considers a construction of state space that begins with the set of pieces of information (“propositions in logic, statements in English or another language, or mathematical formulas”) denoted by Φ . A state is defined as a subset of Φ . Lipman (1999) relaxes the assumption of the agent being logically omniscient, and determines conditions under which the agent’s behavior fits within the SEU theory. My paper differs from Lipman (1999) since, in general, my procedure need not depend on the violation of logical omniscience. Moreover, in contrast to my paper, Lipman (1999) assumes that the researcher knows the set Φ which implies that the researcher knows how the agent constructs her state space.

Although it does not develop a new theory, this paper associates alternatives with subsets of the set of consequences X and, consequently, is related to the literature concerned with an agent who ranks sets of objects. (See Barberà et al. (2004) for the most complete review of this literature.) The models in this literature consist of two stages. In the first stage, the agent chooses a subset of X . What happens at the second stage divides these models into the following two groups. Group A consists of models in which, at the second stage, Nature picks the final consequence (see section 3 in Barberà et al. (2004)). Group B consists of models in which, at the second stage, the agent picks the final consequence (see section 4 in Barberà et al. (2004)).

The models in both of these groups share one feature, which separates them from the setup I propose. They all consider the agent’s binary order to be defined over the subset of the

power set of X , a construction that implicitly assumes that two alternatives with the same set of consequences should be perceived as the same object or, at least, that the agent should be indifferent between them. However, as noted in section 2, this approach is not appropriate for this paper. Thus, in this paper, the agent ranks options, rather than subsets of X .

The models in Group A are concerned with “complete uncertainty” which, according to Barberà et al. (2004), “refers to a situation where the agent knows the set of possible consequences of an action but cannot assign probabilities to those outcomes.” For this reason, these models are not related to this paper. However, one important set of models on complete uncertainty can be interpreted as if they were considering an agent who constructs probabilities over the subsets of X . These are the models which build on the concept of Hurwicz’s α -criterion (see Arrow and Hurwicz (1972) and Hurwicz (1951)) and develop the idea of evaluating sets by taking only their best and worst consequences into consideration.² A subset Y of X is evaluated by $V_\alpha(Y) = \alpha(Y) \cdot u(\max Y) + (1 - \alpha(Y)) \cdot u(\min Y)$, where α is a function from the set of subsets of X to $(0, 1)$, and where $\max Y$ denotes the best element of Y , and $\min Y$ the worst. The function α can be interpreted as a probability on Y . This functional representation resembles the result obtained by Ghirardato (2001). In his model, the state space Ω is fixed, and the agent perceives acts as set-valued functions. In fact, V_α can be considered to be an extreme version of Ghirardato (2001), in which the agent perceives no state space. However, such a representation is not appropriate for the purpose of this paper, since it does not capture the agent’s perception of alternatives. Since Y represents all consequences deemed possible, I focus on the probability measures with support Y . In V_α , however, the support of α is not Y but rather $\{\max Y, \min Y\}$.

For Group B, the relevant research is in the literature on “preference for flexibility” that begins with Kreps (1979) and examines the agent who is not sure about her future preferences. This uncertainty is captured by a state space S in which a state s represents a possible future binary order over set X . My paper implicitly assumes that the agent does not face uncertainty regarding her future self. Consequently, the construction of the state space here differs from the construction proposed in Kreps (1979), in which the agent evaluates a subset Y of X by $V(Y) = \sum_{s \in S} \max_{y \in Y} u(s, y)$, where $u(s, y)$ is an ex-post, state-dependent utility function. In contrast, in my paper a menu is evaluated by its subjective expected utility, and the utility over X is state-independent. Kreps (1992) reinterprets the original “preference for flexibility” model as representing the agent who faces unforeseen contingencies—that is,

²See, for instance, Barberà et al. (1984), Barberà and Pattanaik (1984), Ben Larbi et al. (2010), Bossert (1989), Bossert et al. (1994), Bossert et al. (2000), Dutta and Sen (1996), Kannai and Peleg (1984), Nehring and Puppe (1996), Olszewski (2007), Pattanaik and Xu (1998), and Pattanaik and Xu (2000).

the agent who is unable to account for all possible future uncertainties. In contrast, my paper does not represent the agent who lacks a complete picture of uncertainty, but rather the researcher who does not know which uncertainties the agent has taken into account.

5 Conclusions

If the researcher, who tests the SEU theory, assumes some state space Ω and rejects the SEU theory, then his conclusion is not unconditionally valid. There may exist an alternative state space $\tilde{\Omega}$ and a probability measure $\tilde{\lambda}$ over $\tilde{\Omega}$ which together do not lead to the rejection of the SEU theory. If such $\tilde{\Omega}$ and $\tilde{\lambda}$ exist, then the agent's behavior is SEU-rationalizable. This paper shows how to SEU-rationalize the agent's behavior and when an SEU-rationalization is possible.

I believe that the assumptions which allow to SEU-rationalize the agent's behavior are not demanding. It is difficult to consider violations of Assumption 1—incompleteness and intransitivity—as valid reasons to reject the SEU theory. I leave it for the reader to decide whether possible violations of Assumption 2 constitute convincing or interesting arguments against the SEU theory.

I would like to stress that if the agent's behavior can be SEU-rationalized, then this does not mean that the SEU theory has been *accepted* and non-SEU theories³ rejected. Rather, the correct conclusion would be that the SEU theory has *not been rejected* and that other theories may also be able to explain the agent's behavior by relying on the assumption of a distinct state space. Consequently, instead of saying that my procedure allows the researcher to accept the SEU theory, I say—more precisely—that it allows him to not reject it. For this reason, this paper should not be interpreted as a critique of non-SEU theories. Rather, it is a critique of motivating such theories by only indicating the evidence which rejects the SEU theory.

6 Appendix

I extend my result to infinite sets of acts and consequences. Fix $(\Omega, X, \mathcal{F}, \succsim)$ with \succsim violating the SEU theory. I impose the following assumptions.

³See Amarante (2014), Eichberger and Kelsey (2009), Ghirardato (2010), Gilboa and Marinacci (2011), and Siniscalchi (2008) for a review of the non-SEU literature.

Assumption 3

1. X is a separable metric space.
2. If X_f is uncountable, then X_f is a closed subset of X .
3. \mathcal{F} is a separable metric space.
4. For each $x \in X$, there is $f_x \in \mathcal{F}$.
5. \succsim is a complete, transitive, and continuous binary order over \mathcal{F} .

Assumption 4

1. For each $f \in \mathcal{F}$, there exist $f_{\max f}$ and $f_{\min f}$.
2. For each $f \in \mathcal{F}$, if $f_{\max f} \succ f_{\min f}$, then $f_{\max f} \succ f \succ f_{\min f}$.

Remark: Instead of assuming the existence of $f_{\max f}$ and $f_{\min f}$, one can require that X be a compact metric space. In that case, X_f , a closed subset of X , is also compact, and since u is a continuous function, the Weierstrass Theorem dictates that both $f_{\max f}$ and $f_{\min f}$ exist.

Proposition 2

Fix a quadruple $(\Omega, X, \mathcal{F}, \succsim)$. For each $f \in \mathcal{F}$, fix a set X_f . Assume the Axiom of Choice. If Assumptions 3 and 4 hold, then an SEU-rationalization of \succsim is possible.

In the first stage of SEU-rationalization, it is necessary to generate the set of options $\hat{\mathcal{F}}$ and the binary order $\hat{\succsim}$ on $\hat{\mathcal{F}}$. Recall that an option \hat{f} is defined as a pair (f, X_f) where X_f is a subset of X interpreted as the collection of all possible consequences associated with act f (i.e., $X_f = f(\Omega)$). Let $\hat{\lambda}_f$ be a probability measure associated with option \hat{f} . It is important that this probability measure captures the agent's perception of \hat{f} . In the case of a countable X_f , if x is considered to be possible when f is chosen, then $\hat{\lambda}_f$ clearly assigns a nonzero weight to x . However, with uncountable X_f , it is not possible to require that $\hat{\lambda}_f(x) > 0$ for each $x \in X_f$. This implies that there is no straightforward definition of the "possible consequences" for uncountable X_f s. Suppose that $X = [0, 1]$ and $X_f = [0, 0.5]$. If $\hat{\lambda}_f(0) = 0$, then it is not clear whether X_f should be defined as $[0, 0.5]$ or $(0, 0.5]$. I solve this problem by implicitly assuming that the agent is indifferent between X_f and the closure of X_f . Thus, in Assumption 3, I focus only on those uncountable X_f s that are closed in X .

Consequently, $\hat{\lambda}_f$ is said to capture the agent's perception about an act whenever the support of $\hat{\lambda}_f$ is X_f . (Note that if λ is a probability measure defined on a metric space X , then the support of λ is defined as the smallest closed subset C of X such that $\lambda(C) = 1$.) However, without additional requirements, it is not necessarily true that, for a given X_f , there exists a probability measure with support X_f . This issue is also resolved by adopting Assumption 3, which requires that an uncountable X is a separable metric space. This guarantees that, for a given closed subset of X , X_f , there is a probability measure on X with support X_f (see Exercise 4 in Chapter 7.1 in Dudley (2002)).

Assumption 3 guarantees the existence of \hat{V} (see Debreu (1954, 1964)), with u defined as $u(x) := \hat{V}(\hat{f}_x)$. It remains to be proven that (a) for each option there exists a probability measure $\hat{\lambda}_f$ with support X_f , and (b) that $\hat{V}(\hat{f}) = \int_X u(x)d\hat{\lambda}_f$.

Let $P(X)$ denote the collection of probability measures on X that are Borel for uncountable X . Fix \hat{f} and let $P_f(X)$ be a subset of $P(X)$ that consists of probability measures on X with support X_f . For countable X_f , the nonemptiness of $P_f(X)$ is obvious. For uncountable X_f , the fact that X_f is a closed subset of a metric separable space (Assumption 3) guarantees the nonemptiness of $P_f(X)$. Let $\Psi : P(X) \rightarrow \mathbb{R}$ be the integral function $\Psi(\lambda) := \int_X u d\lambda$. The remaining part of proof shows that, for a given \hat{f} , there exists $\hat{\lambda}_f \in P_f(X)$ such that $\Psi(\hat{\lambda}_f) = \hat{V}(\hat{f})$.

Consider an option \hat{f} such that $\hat{f}_{\max f} \succ \hat{f}_{\min f}$. The idea of the proof is to find two probability measures, λ_1 and λ_2 , that both belong to $P_f(X)$ such that $c_1 = \Psi(\lambda_1) < \hat{V}(\hat{f}) < \Psi(\lambda_2) = c_2$. With c_1 and c_2 in hand, it is possible to find the unique α such that $\alpha c_1 + (1 - \alpha)c_2 = \hat{V}(\hat{f})$. With this α , the desired $\hat{\lambda}_f$ is defined as $\hat{\lambda}_f = \alpha\lambda_1 + (1 - \alpha)\lambda_2$.

It remains to be proven that, for a given \hat{f} , such probability measures, λ_1 and λ_2 , indeed exist. Note that due to Assumption 4, $\Psi(\delta_{\min f}) < \hat{V}(\hat{f}) < \Psi(\delta_{\max f})$, where δ_x is a degenerate probability measure with mass 1 at x . One way to show that there always is a λ_1 such that $\Psi(\lambda_1) < \hat{V}(\hat{f})$ involves proving that there is a sequence of probability measures $\{\lambda_1^n\}$ such that each $\lambda_1^n \in P_f(X)$ and $\Psi(\lambda_1^n) \rightarrow \Psi(\delta_{\min f})$. The existence of such a sequence guarantees that there is some N such that $\Psi(\delta_{\min f}) < \Psi(\lambda_1^N) < \hat{V}(\hat{f})$. Such a λ_1^N then becomes the desired λ_1 . Since the existence of λ_2 is proven in the same way, I focus on λ_1 .

Consider a countable X_f . I construct $\{\lambda_1^n\}$ in the following way. Let $\lambda_1^n(\min f) = 1 - \frac{1}{n}$ and let $\sum_{x \in X_f \setminus \{\min f\}} \lambda_1^n(x) = \frac{1}{n}$. If X_f is finite, with cardinality K , then let $\lambda_1^n(x) = \frac{1}{(K-1)n}$ for $x \in X_f \setminus \{\min f\}$. If X_f is infinite, then let $X_f = \{x, x_0, x_1, \dots\}$, where $x = \min f$. For x_m , define $\lambda_1^n(x_m) = \frac{1}{2^n} \cdot \frac{1}{2^m}$. Note that $\sum_{m=0}^{\infty} \lambda_1^n(x_m)$ is a convergent geometric series such that $\sum_{m=0}^{\infty} \frac{1}{2^n} \cdot \frac{1}{2^m} = \frac{1}{n}$. With $\{\lambda_1^n\}$ determined, it is then possible to compute $\Psi(\lambda_1^n) =$

$\sum_{x \in X_f} u(x) \lambda_1^n(x)$; observe that the sum converges, as desired, to $\Psi(\delta_{\min f})$ as n goes to infinity. Finally, consider an uncountable X_f . As already noted, the fact that X is separable and X_f is a closed subset of X implies the existence of probability measure $\hat{\lambda}_f$ on X , with the support X_f . Let λ_1^n be a weighted measure that assigns $\frac{1}{n}$ to $\hat{\lambda}_f$ and $1 - \frac{1}{n}$ to $\delta_{\min f}$. Each λ_1^n then has support X_f , and $\Psi(\lambda_1^n) \rightarrow \Psi(\delta_{\min f})$ when $n \rightarrow \infty$.

By the Axiom of Choice, for an arbitrary $\hat{\mathcal{F}}$, there exists $\tilde{\Omega}$ defined as $\tilde{\Omega} := \times_{f \in \mathcal{F}} X_f$. Let τ be a finite nonempty subset of \mathcal{F} . Let $\tilde{\Omega}_\tau := \times_{f \in \tau} X_f$. Let $\tilde{\lambda}_\tau$ be the probability measure on $\tilde{\Omega}_\tau$, constructed as a product of the probability measures $\hat{\lambda}_f$ s. By Proposition V.1.2 in Neveu (1965), there exists the unique probability measure $\tilde{\lambda}$ on $\tilde{\Omega}$, which agrees with each $\tilde{\lambda}_\tau$ on cylinders. An option \hat{f} is transformed into act \tilde{f} by defining $\tilde{f} : \tilde{\Omega} \rightarrow X$ as a projection function. Consequently, function \tilde{f} is measurable. The binary order $\tilde{\succsim}$ over $\tilde{\mathcal{F}}$ is as $\tilde{f} \tilde{\succsim} \tilde{g}$ if and only if $\hat{f} \hat{\succsim} \hat{g}$. Finally, define $\tilde{V} : \tilde{\mathcal{F}} \rightarrow \mathbb{R}$ as $\tilde{V}(\tilde{f}) := \int_{\tilde{\Omega}} u(\tilde{f}(\tilde{\omega})) d\tilde{\lambda}$.

Since $\tilde{\lambda}$ is a product measure and \tilde{f} is a projection function, they induce a probability measure on X that is the same as $\hat{\lambda}_f$ (condition 2 in Definition 2.1). Note that $\tilde{\succsim}$ is constructed in a consistent way from $\hat{\succsim}$ (condition 3 in Definition 2.1). By the construction of $\tilde{\lambda}$ and \tilde{f} , it follows that $\tilde{V}(\tilde{f}) = \hat{V}(\hat{f})$. This implies that $\tilde{\succsim}$ has an SEU representation (condition 4 in Definition 2.1). This ends the proof of Proposition 2.

References

- AMARANTE, M. (2014): “What is Ambiguity?” *CIREQ Cahier 04-2014*.
- ARROW, K. J. AND L. HURWICZ (1972): “An Optimality Criterion for Decision-Making under Ignorance,” in *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*, ed. by C. Carter and J. Ford, Oxford, UK: Basil Blackwell, 1–11.
- BARBERÀ, S. (1977): “The Manipulation of Social Choice Mechanisms that Do Not Leave “Too Much” to Chance,” *Econometrica*, 45, 1573–1588.
- BARBERÀ, S., C. R. BARRETT, AND P. K. PATTANAIK (1984): “On Some Axioms for Ranking Sets of Alternatives,” *Journal of Economic Theory*, 33, 301–308.
- BARBERÀ, S., W. BOSSERT, AND P. K. PATTANAIK (2004): “Ranking Sets of Objects,” in *Handbook of Utility Theory*, ed. by P. Hammond and C. Seidl, Boston, MA: Kluwer Academic Publishers, vol. 2: Extensions, chap. 17, 893–978.

- BARBERÀ, S. AND P. K. PATTANAIK (1984): “Extending an Order on a Set to the Power Set: Some Remarks on Kannai and Peleg’s Approach,” *Journal of Economic Theory*, 32, 185–191.
- BEN LARBI, R., S. KONIECZNY, AND P. MARQUIS (2010): “A Characterization of Optimality Criteria for Decision Making under Complete Ignorance,” *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning, Toronto*, 172–181.
- BILLOT, A. AND V. VERGOPOULOS (2014): “Dynamic Consistency and Expected Utility with State Ambiguity,” *PSE Working Papers 2014-19*.
- BLUME, L., D. A. EASLEY, AND J. Y. HALPERN (2009): “Constructive Decision Theory,” *Economics Series, Institute for Advanced Studies, Vienna*, No. 246.
- BLUME, L. E., D. A. EASLEY, AND J. Y. HALPERN (2006): “Redoing the Foundations of Decision Theory,” *Tenth International Conference on Principles of Knowledge Representation and Reasoning*, 10–24.
- BOSSERT, W. (1989): “On the Extension of Preferences over a Set to the Power: An Axiomatic Characterization of a Quasi-ordering,” *Journal of Economic Theory*, 49, 84–92.
- BOSSERT, W., P. K. PATTANAIK, AND Y. XU (1994): “Ranking Opportunity Sets: An Axiomatic Approach,” *Journal of Economic Theory*, 63, 326–345.
- (2000): “Choice under Complete Uncertainty: Axiomatic Characterizations of Some Decision Rules,” *Economic Theory*, 16, 295–312.
- DEBREU, G. (1954): “Representation of a Preference Ordering by a Numerical Function,” in *Decision Processes*, ed. by R. Thrall, C. Coombs, and R. Davis, New York, NY: Wiley, 159–165.
- (1964): “Continuity Properties of Paretian Utility,” *International Economic Review*, 5, 285–293.
- DUDLEY, R. M. (2002): *Real Analysis and Probability*, New York, NY: Cambridge University Press.
- DUTTA, B. AND A. SEN (1996): “Ranking Opportunity Sets and Arrow Impossibility Theorems: Correspondence Results,” *Journal of Economic Theory*, 71, 90–101.

- EICHBERGER, J. AND D. KELSEY (2009): “Ambiguity,” in *The Oxford Handbook of Rational and Social Choice*, ed. by P. Anand, P. Pattanaik, and C. Puppe, Oxford: Oxford University Press, 113–139.
- EPSTEIN, L. G. (2010): “A Paradox for the “Smooth Ambiguity” Model of Preference,” *Theoretical Economics*, 78, 2085–2099.
- GHIRARDATO, P. (2001): “Coping with Ignorance: Unforeseen Contingencies and Non-additive Uncertainty,” *Economic Theory*, 17, 247–276.
- (2010): “Ambiguity,” in *Encyclopaedia of Quantitative Finance*, ed. by R. Cont, Chichester, UK: J. Wiley and Sons, 39–44.
- GILBOA, I. (2009): *Theory of Decision under Uncertainty*, New York, NY: Cambridge University Press.
- GILBOA, I. AND M. MARINACCI (2011): “Ambiguity and the Bayesian Paradigm,” in *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, New York: Cambridge University Press.
- GILBOA, I. AND L. SAMUELSON (2012): “Subjectivity in Inductive Inference,” *Theoretical Economics*, 7, 183–215.
- GILBOA, I. AND D. SCHMEIDLER (1994): “Additive Representations of Non-additive Measures and the Choquet Integral,” *Annals of Operations Research*, 52, 43–65.
- (2004): “Subjective Distributions,” *Theory and Decision*, 56, 345–357.
- GRANT, S., J. KLINE, I. MENEGHEL, J. QUIGGIN, AND R. TOURKY (2015): “A Theory of Robust Experiments for Choice under Uncertainty,” *working paper*.
- HURWICZ, L. (1951): “Optimality Criteria for Decision Making Under Ignorance,” *Cowles Commission Discussion Paper: Statistics*, No. 370.
- KADANE, J. B. (1992): “Healthy Scepticism as an Expected-utility Explanation of the Phenomena of Allais and Ellsberg,” *Theory and Decision*, 32, 57–64.
- KANNAI, Y. AND B. PELEG (1984): “A Note on the Extension of an Order on a Set to the Power Set,” *Journal of Economic Theory*, 32, 172–175.
- KARNI, E. (2008): “Unknowable States and Choice-based Definitions of Subjective Probabilities,” *Economics Letters*, 99, 534–536.

- KARNI, E. AND D. SCHMEIDLER (1991): “Utility Theory with Uncertainty,” in *Handbook of Mathematical Economics*, ed. by W. Hildenbrand and H. Sonnenschein, Amsterdam: Elsevier, vol. 4, 1763–1831.
- KARNI, E. AND M.-L. VIERØ (2013): “Reverse Bayesianism: A Choice-Based Theory of Growing Awareness,” *American Economic Review*, 103, 2790–2810.
- KREPS, D. M. (1979): “A Representation Theorem for “Preference for Flexibility”,” *Econometrica*, 47, 565–577.
- (1992): “Static Choice and Unforeseen Contingencies,” in *Economics Analysis of Markets and Games: Essays in Honor of Frank Hahn*, ed. by P. Dasgupta, D. Gale, O. Hart, and E. Maskin, Cambridge, MA: The MIT Press, 258–281.
- LEFEVRE, J.-A., M. PENNER-WILGER, A. A. PYKE, T. SHANAHAN, AND W. A. DESLAURIERS (2014): “Putting Two and Two Together: Declines in Arithmetic Fluency among Young Canadian Adults, 1993 to 2005,” *Carleton University Cognitive Science Technical Report 2014-01*.
- LIPMAN, B. L. (1999): “Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality,” *Review of Economic Studies*, 66, 339–361.
- MACHINA, M. J. (2003): “States of the World and the State of Decision Theory,” in *The Economics of Risk*, ed. by D. Meyer, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 17–49.
- NEHRING, K. AND C. PUPPE (1996): “Continuous Extension of an Order on a Set to the Power Set,” *Journal of Economic Theory*, 68, 456–479.
- NEVEU, J. (1965): *Mathematical Foundations of the Calculus of Probability*, San Francisco, CA: Holden-Day.
- OLSZEWSKI, W. (2007): “Preferences Over Sets of Lotteries,” *Review of Economic Studies*, 74, 567–595.
- PATTANAIK, P. K. AND Y. XU (1998): “On Preference and Freedom,” *Theory and Decision*, 44, 173–198.
- (2000): “On Ranking Opportunity Sets in Economic Environments,” *Journal of Economic Theory*, 93, 48–71.

- SAVAGE, L. J. (1972): *The Foundations of Statistics*, New York, NY: Dover Publications (revised and expanded version of a work originally published by John Wiley & Sons in 1954).
- SCHIPPER, B. C. (2013): “Awareness-Dependent Subjective Expected Utility,” *International Journal of Game Theory*, 42, 725–753.
- SCHMEIDLER, D. (1989): “Subjective Probability and Expected Utility Without Additivity,” *Econometrica*, 57, 571–587.
- SCHMEIDLER, D. AND P. WAKKER (1987): “Expected Utility and Mathematical Expectation,” in *The New Palgrave Dictionary of Economics*, ed. by J. Eatwell, M. Milgate, and P. Newman, New York, NY: Palgrave Macmillan, first ed.
- SINISCALCHI, M. (2008): “Ambiguity and Ambiguity Aversion,” in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, New York, NY: Palgrave Macmillan, second ed.
- STANDING, L. G. (2006): “Why Johnny Still Can’t Add: Predictors of University Students’ Performance on an Elementary Arithmetic Test,” *Social Behavior and Personality: An International Journal*, 34, 151–160.
- STANDING, L. G., R. A. SPROULE, AND A. LEUNG (2006): “Can Business and Economics Students Perform Elementary Arithmetic?” *Psychological Reports*, 98, 549–555.