Decisions and Discovery: A Non-Bayesian Perspective

Edi Karni [*] [†]
Johns Hopkins University

May 31, 2025

Abstract

The paper explores a new, non-Bayesian, approach to modeling choice behavior of decision maker's who is aware that there may exist potential outcomes of his actions of which he is not aware. The new approach suggests an adaptation of Ewens (1972) generalization of De Morgan's (1838) formula of the probability of known and unpredictable outcomes to the context of decision making under risk, and embeds it in an expected utility model with costly actions.

Keywords: Awareness, unawareness, inductive inference, predicting the unpredictable.

JEL classification: D8, D81, D83

[*] Johns Hopkins University, Department of Economics. E-mail: karni@jhu.edu.

[†] I am grateful to Burkhard Schipper for his comments and suggestions.

1 Introduction

Decisions are choices among feasible courses of action the consequences of which have been discovered over time. In the case of decision making under uncertainty, decision makers' beliefs about the likely occurrence of the consequences are formed by inductive inference. It is often the case that decision makers encounter a consequence of whose existence they were unaware (i.e., a consequence that, before it was obtianed for the first time, was unknown unknown). Repeated discovery of unanticipated, novel, consequences, however, alerts decision makers to the possibility that there may exist additional, unanticipated, consequences whose nature they are unable to conceive of. In other words, decision makers may be aware of their unawareness, and this awareness of unawareness may affect their choice behavior. Invoking the Bayesian approach, Karni and Vierø(2017) proposed a model of decision making under uncertainty that incorporates the decision makers' awareness of unawareness.

In this paper I explore a different, non-Bayesian, approach to modeling decision making under uncertainty, based on inductive inference, that incorporates the notion of awareness of unawareness. The model assigns objective probabilities and utilities to unanticipated, inconceivable, consequences. To motivate this exploration, I consider two instances requiring making decision under uncertainty, in which decision makers have access to data that may be used to calibrate the likelihoods of the known possible consequences of their actions, and quantify the probability of encountering unanticipated consequences.

Example 1 - A decision maker who must choose whether or not to vaccinate against a disease and, if the decision is to vaccinate, which of several available vaccines to take. Clinical trials in which different vaccines are tested and, once approved and implemented, the cumulative evidence regarding their effectiveness and potential side effects provide the data on which the decision makers may base decisions. Because the testing and using new vaccines is a process of exploring uncharted terrain, the potential of discovering novel, previously unsuspected, health consequences is ubiquitous. In deciding whether to vaccinate, or which vaccine to choose, decision makers use the accumulated evidence to form beliefs about the likelihoods of occurrence of known outcomes and the potential existence of unforeseeable health effects.

Example 2- For decades statins have been used to reduce cholesterol levels. Their side effects, are well known, leaving little chance to discovering new, unanticipated, side effects, both in the short and the long run.[1] Clinical trials have shown that a new medication is more effective in lowering the cholesterol levels and has fewer undesirable side effects. However, the lack of experience with the new medication implies that there is greater chance of discovering unanticipated side effects in the long run. A decision maker who must choose between the familiar statins and the less familiar new medication may take into account the likelihoods of the known effects as well as the potential of unforeseeable health consequences.

These examples illustrate the need for theories of decision making under uncertainty, founded on inductive inference, that accommodate the potential existence of unanticipated consequences. With few recent exceptions, all the theories of decision making under uncertainty maintain that the set of the ultimate outcomes, or payoffs, are known. To the extent that there is learning, it is expressed as the updating of subjective probabilities on a fixed state space. The exceptions include recent models of decision making under uncertainty in which the decision makers are not assumed to be aware of all the possible consequences that may result from their choice of actions, and may also be aware of their unawareness. Invoking the Bayesian approach, Karni and Vierø (2013), addressed this by expanding the state space and axiomatized a process, dubbed 'reverse Bayesianism', according to which the decision maker's updates her beliefs following a procedure that maintain the spirit of Bayes' rule. This approach was further explored and elaborated in Karni and Vierø (2015, 2017), Dominiak and Tserenjigmid (2018), Karni, Vierø, Valenzuela-Stookey (2021), Chakravarty, Kelsey, and Teitelbaum (2021), Vierø (2021)[2].

A non-Bayesian, approach to modeling the process of exploration and discovery in an environment in which unsuspected events may occur has been pursued in probability theory. The problem is what to do when such events obtain. In other words,

"How can we predict the occurrence of something we neither know, nor

[1] The known side effects include headache, dizziness, feeling unusually tired or physically weak, digestive system problems, such as constipation, diarrhoea, or indigestion, muscle pain, sleep problems, low blood platelet count.

[2] The study unawareness is also taken up in epistemologic game theory by Heifetz, Meier, and Schipper, (2006), (2008), (2013) and Grant and Quiggin (2013). For experimental test of reverse Bayesianism see Becker, Melkonyan, Proto, Sofianos, and Trautman, (2020).

even suspect, exists? Subjective probability and Bayesian inference, despite their many impressive successes, would seem at a loss to handle such a problem given their structure and content." [Zabell (1992) p. 206].

A particular instance of this difficulty is the so-called *sampling of species problem.*[3] The process may best be described as follows:

> "Imagine that we are in a new terrain, and observe the different species present. Based on our past experience, we may anticipate seeing certain old friends - black crows, for example - but stumbling across a giant panda may be a complete surprise. And, yet, all such information will be grist to our mill: if the region is found rich in the variety of species present, the chance of seeing a particular species again may be judged small, while if there are only a few present, the chances of another sighting will be judged quite high. The unanticipated has its uses." [Zabell (1992) p. 206]

De Morgan (1838) proposed an updating process for dealing with precisely this issue. According to De Morgan, if following a sequence, $Z_1, Z_2, ..., Z_N$ of N trials (i.e., observations) t categories, or outcomes, labeled $c_1, ..., c_t$, have been observed, then the probability of seeing the outcome on trial $N+1$ fall into the $j-th$ category is:

$$\Pr\{Z_{(N+1)} = c_j\} = \frac{n_j + 1}{N + t + 1}, \tag{1}$$

where n_j denotes the number of times each of the t outcomes occurred in N trials. Notice that $\Sigma_{j=1}^{t} \Pr\{Z_{(N+1)} = c_j\} < 1$ (i.e., the probability of observing, in the $N+1$ trial an outcome seen before is smaller than 1). This formula implicitly assigns a category not yet observed, denoted $\hat{c}$, a probability of occurring equal to

$$\Pr\{Z_{(N+1)} = \hat{c}\} = \frac{1}{N + t + 1}. \tag{2}$$

The stochastic process depicted by De Morgan's proposal is generated by the following urn model.[4] Consider an urn containing t balls of different colors and a black ball called the *mutator*. Draw a ball at random. If a colored ball is drawn, then it is replaced and

[3] Zabell (1992) provides an insightful discussion and numerous references.

[4] See Zabell (1992).

another ball of the *same* color is added. If the mutator is drawn, then it is replaced and another ball of *new* color is added.

Let n_j be the number of times a ball of color c_j is drawn in N trials and let $\mathbf{n} = (n_1, ..., n_t)$ be the frequency distribution of the known colors. What is the probability distribution of the next draw from the same urn? The answer to this question is given by the De Morgan formulas (1) and (2).[5] Ewens (1972) generalized De Morgan's formula allowing the mutator assume weight different from one.

The sole concern of the De Morgan process described above is epistemic - the exploration of the same 'new terrain' through repeated observations using the same procedure (e.g., sampling from the same urn). In particular, it is not concerned with the possibility of exploration using alternative procedures (e.g., sampling from different urns) Furthermore, the observations, species or colors, having no intrinsic welfare implications, are purely informative.

In this paper, I explore the application of the Ewens' process to situations in which decision makers facing choice under uncertainty are aware of the possibility that there exist unanticipated outcomes. This objective requires the modification and extension of stochastic process described above. Specifically, we have to consider repeated observations of outcomes generated by the choice of alternative courses of action, (e.g., sampling from different urns) while taking into account that information acquired under one course of action informs the decision maker about the possibility of the occurrence and prevalence of outcomes under other courses of action. In particular, discovering an outcome never seen before informs the decision maker of its existence, thereby making him aware of the possible occurrence of this outcome under all courses of action. The correlations among the samples from distinct urns necessitate the loss of the property of partition exchangeability that characterizes the prediction rule of De Morgan (1838) and its generalization due to Ewens (1972). Furthermore, in addition to exploration, depicted by the sampling of species problem, the choice of alternative courses of action involves exploitation – the outcomes have material (i.e., welfare) consequences – and may involve distinct direct or indirect costs.

The next section describes the extension of the Ewens' (1972) generalization of the De

[5] At a deeper level, the De Morgan formula is generated by the representation of random partitions exchangeability. A more detailed exposition of this idea is beyond the scope of this paper. The interested reader will find an excellent review and references in Zabell (1992).

Morgan proposal. Section 3 introduces a decision making model. Section 4 includes a discussion of the exploration-exploration tradeoff implied be the model, its relation to the multi-armed bandit models, and a brief review of the related literature.

2 Generalized Sampling Process

2.1 Similarity

The extended sampling process is best described by the following multi-urn model that allows the mutator to assume different weights (e.g., different numbers of mutators in the different urns) and the random draws from distinct urns to be stochastically dependent.

Let $\mathcal{U} := \{U_1, ..., U_m\}$ be a finite set of *urns*. Let N_i be the number of draws from U_i, $i = 1, ..., m$, containing balls of t_i different colors, and some black balls. Denote by θ_i the weight of the mutator (i.e., black balls) in urn U_i.

Consider the following process. Select an urn from $\mathcal{U}$ and draw a ball at random from the selected urn. If a colored ball is drawn, then it is replaced and another ball of the same color is added. If a black ball is drawn, it is replaced and another ball of new color is added.

Define $C_i = \{c_{i1}, ..., c_{it_i}\}$, $i = 1, ..., m$, the set of colors, or categories, observed in N_i draws from U_i. Let $C = \cup_{i=1}^m C_i$ be the set of colors known to exist after $N = \Sigma_{i=1}^m N_i$ draws form all the urns. Let n_{ic_k} be the number of balls of color $c_k \in C$ observed in N_i draws from U_i. Denote by $\mathbf{n}_{ic}(C) = \left(n_{ic_1}, ..., n_{ic_{|C|}}\right)$ the frequency of draws of the known colors after a sequence $Z_{i1}, Z_{i2}, ..., Z_{iN_i}$ of draws from U_i. Note that $n_{ic_k} = 0$ if $c_k \in C \backslash C_i$. In other words, $c_k \in C \backslash C_i$ is known to exist but has not been observed in U_i.

Given $N_1, ..., N_m$ and C, let $p_{ic_k} := n_{ic_k} / (N_i + \theta_i)$, $k \in \{1, ..., \mid C \mid\}$, and $p_{i\theta_i} := \theta_i / (N_i + \theta_i)$. By definition, $p_{ic_k} = 0$ if $c_k \in C \backslash C_i$. Let $\mathbf{p}_i := (p_{ic_1}, ..., p_{ic_{|C|}}, p_{i\theta_i})$, $i = 1, ..., m$. For every pair of probability vectors $\mathbf{p}_i, \mathbf{p}_j$, $i, j \in \{1, ..., m\}$ denote by $\langle \mathbf{p}_i, \mathbf{p}_j \rangle$ their inner product, and let

$$\kappa(i, j) := \frac{\langle \mathbf{p}_i, \mathbf{p}_j \rangle}{\| \mathbf{p}_i \| \| \mathbf{p}_j \|}. \quad (3)$$

Then, $\kappa(i, j) = \cos \tau(i, j)$ is a measure of the angle between the two probability vectors $\mathbf{p}_i$ and $\mathbf{p}_j$. I interpret $\kappa(i, j)$ to be a measure of *similarity* between the underlying stochastic processes that depict the draws from N_i and N_j.

Because the probability vectors are in $\mathbb{R}_+^n$, the similarity measure takes values in the unit interval. Obviously, $\kappa(i,i) = 1$, for all $i = 1, ..., m$, (i.e., each urn is perfectly similar to itself) and $\kappa(i,j) = 0$ if and only if $\mathbf{p}_i$ and $\mathbf{p}_j$ are orthogonal. Note, however, that under the assumption that all the urns contain mutators, no two probability vectors $\mathbf{p}_i$ and $\mathbf{p}_j$ are orthogonal. Hence, $\kappa(i,j) > 0$, and $\lim_{N_i \to \infty} \kappa(i,j) = 0$, for all $i, j = 1, ..., m$. The larger values of $\kappa(i,j)$, the more similar are the content of the urns U_i and U_j. It is worth underscoring that, before the sampling starts, all the urns are perfectly similar (i.e., $\kappa(i,i) = 1$, for all $i, j = 1, ..., m$), in the sense that nothing is known about their content.

The similarity measure is monotonic decreasing in the difference between the size of the samples of the urns. To grasp this, consider two urns, say U_i and U_j, and suppose that the corresponding number of draws are N_i and N_j, where $N_j > N_i$. Suppose further that, given N_i and N_j, the conditional probabilities, $\hat{\mathbf{p}}_i := \mathbf{p}_i / (1 - p_{i\theta_i})$ and $\hat{\mathbf{p}}_j := \mathbf{p}_j / \left(1 - p_{j\theta_j}\right)$ are the same. However, the larger is the difference in the sample sized (i.e., the larger is the difference between the weights of evidence), the larger is the difference $p_{j\theta_j} - p_{i\theta_i}$. Consequently, the smaller is the similarity coefficient $\kappa(i,j)$. In general, two urns are more similar when high-probability categories in one urn are high-probability categories in the other urn and low-probability categories in one urn are low-probability categories in the other urn.

Given $\mathbf{p}_i$, $i = 1, ...m$, the probability of observing $c \in C$ conditional on a draw from U_i is:

$$\Pr\{Z_{i(N_i+1)} = c \mid \mathbf{p}_1, ..., \mathbf{p}_m\} = \sum_{j=1}^{m} p_{jc} \kappa(i,j) \frac{N_j}{N}. \tag{4}$$

The probability of encountering a color not seen before, (i.e., the probability of drawing θ_i) is

$$\Pr\{Z_{i(N_i+1)} = \theta_i \mid \mathbf{p}_1, ..., \mathbf{p}_m\} = \sum_{c \in C} (1 - \sum_{j=1}^{m} p_{jc} \kappa(i,j) \frac{N_j}{N}). \tag{5}$$

In the analysis that follows, I assume the decision maker predicts the outcomes of his actions using these formulas. If $c \in C \backslash C_i$ then it is known to exist, and is assigned probability on the basis of the similarity of U_i and the urns in which c was observed, even though it didn't show up in the sampling from U_i. The decision maker is unaware of outcomes that are not in C.[6]

[6] De Morgan's and Ewen's prediction rules are charaterized by partition exchangeability. The prediction

2.2 State-space formulation

A cornerstone of Bayesian theories of decision making under uncertainty is a primitive, immutable, state space, whose elements represent complete resolutions of uncertainty. Specifically, the presumption is that there is a *true state*, that is "a description of the world so complete that, if true and known, the consequences of every action would be known." (Arrow [1965]).

The decision model of this paper does not invoke the notion of states as a primitive concept. It is possible, however, to derive, within the framework of the model of this paper, a concept of *evolving state space*. To grasp this, let $A := \{a_1, ..., a_m\}$ denote a (finite) set of courses of actions, and denote by C the set of known consequences, after N choices of actions from A. The state space describing the uncertainty before the $N+1$ choice of action consists of all the mappings from the set of actions to the set of known and unknown outcomes (i.e., $S = \{s : A \to \mathfrak{C}\}$).[7] Clearly, this definition of states represents the resolution of the uncertainty the decision maker faces before choosing the next action. Moreover, given $N_1, ..., N_m$, $\mathbf{p} = (\mathbf{p}_1, ..., \mathbf{p}_m)$ and C, the probability of the state $s = (s(a_1), ..., s(a_m))$ is:

$$\Pr\{s\} = \Pr\{Z_{1(N+1)} = s(a_1) \mid \mathbf{p}\} \times ... \times \Pr\{Z_{m(N+1)} = s(a_m) \mid \mathbf{p}\}.$$

Obviously, the state space in this model is neither primitive nor immutable. In fact, once a new consequence is discovered, the domain of definition of states expands and new states are generated. More concretely, let $\hat{c}$ denote the newly observed outcome and let $C' = C \cup \{\hat{c}\}$ and $\mathfrak{C}' = C' \cup \{\theta\}$, then the new state space is: $S' = \{s' : A \to \mathfrak{C}'\}$. The probabilities of the states are:

$$\Pr\{s'\} = \Pr\{Z_{1(N+2)} = \{s'(a_1) \mid \hat{\mathbf{p}}\} \times ... \times \Pr\{Z_{m(N+2)} = s'\ (a_m) \mid \hat{\mathbf{p}}\},$$

rules (4) and (5) do not satisfy partition exchangeability. To be exact, as long as the sampling is from a single urn, partition exchangeability may be assumed to hold and the prediction rules apply to continuation of draws from the same urn. However, once sampling from another urn is undertaken, cross inferences implies that partition exchangeability no longer applies to either of the urns. This is analogous to the predictive rule of heads and tails in repeated flipping of the same coin, that may be characterized by exchangeablity, whereas the predictive rule corresponding to flipping distinct, correlated, coins do not satisfy exchangeability.

[7] This notion of states was describe in Schmeidler and Wakker (1987) and Karni and Schmeidler (1991) and was invoked in Karni and Vierø(2013, 2017).

where $\hat{\mathbf{p}} = (\hat{\mathbf{p}}_1, ..., \hat{\mathbf{p}}_m)$ and $\hat{\mathbf{p}}_i = (\hat{p}_{i1}, ..., \hat{p}_{i|C'|}, \hat{p}_{i\theta_i}), i = 1, ..., m$, are the updated probabilities vector following the discovery of the new outcome.

If the action a_i that was chosen at the $N+1$ stage results in a known outcome, $c_k \in C$, then the state space does not change but the probabilities of the states do. In particular, let $\mathbf{p}' = (\mathbf{p}'_1, ..., \mathbf{p}'_m)$, where $\mathbf{p}'_i$, $i = 1, ..., m$, are given by (15) then

$$\Pr\{s\} = \Pr\{Z_{1(N+1)} = s(a_1) \mid \mathbf{p}'\} \times ... \times \Pr\{Z_{m(N+1)} = s(a_m) \mid \mathbf{p}'\}.$$

3 The Decision Model

Unlike the Bayesian decision models, in which the state probabilities are subjective, according to the approach of this paper, the probability distributions on the evolving state spaces, induced by the relative frequencies of the outcomes, are objective. Define $\mathfrak{C} = C \cup \{\theta\}$ as the set of *outcomes,* where θ signifies the existence of outcomes not in C. In other words, θ symbolizes unanticipated outcomes whose nature is, by definition, unknown. Hence, *acts correspond to lotteries over the evolving sets of consequences* $\{\mathfrak{C}, \mathfrak{C}', ...\}$. Thus, the decision model depicts choice under risk rather than choice under uncertainty.

The set of acts is finite. Consequently, at any given time (i.e., after any set $N_1, ..., N_m$ of actions), the cumulative data consisting of action-outcome pairs, give rise to a finite set of lotteries. Since no restrictions are imposed on the data,in principle, therefore, the decision maker must consider encountering, and evaluating, any lottery in the space $\Delta(\mathfrak{C})$ of distributions over $\mathfrak{C}$.

3.1 The choice set and the structure of the preference relations

Let $A := \{a_1, ..., a_m\}$ be a set whose elements are alternatives courses of action, or *actions,* for short. The reader may find it convenient to think of $a_i \in A$ as corresponding to $U_i \in \mathcal{U}$. Let $\Delta(\mathfrak{C})$ denote the set of probability distributions on $\mathfrak{C}$.

The product set $\mathbb{C} = A \times \Delta(\mathfrak{C})$ is said to be the *choice set.* A binary relation $\succcurlyeq$ on $\mathbb{C}$, is a *preference relation.* Let $\succ$ and $\sim$ denote the asymmetric and symmetric parts of $\succcurlyeq$, respectively. I assumed that the decision maker is able to rank all pairs, $(a, p) \in \mathbb{C} = A \times \Delta(\mathfrak{C})$. In particular, the decision maker is able to rank such pairs in which the probabilities in $\Delta(\mathfrak{C})$ are obtained from the formulas (4) and (5). For every given $\mathbb{C} = A \times \Delta(\mathfrak{C})$, the structure of the preference relation is depicted axiomatically as follows.

(A.1) **Weak Order -** $\succcurlyeq$ is complete and transitive.

(A.2) **Conditional Archimedean -** For each $a \in A$ and $(a,p), (a,q), (a,r) \in \mathbb{C}$ such that $(a,p) \succ (a,q) \succ (a,r)$ there are $\alpha, \beta \in (0,1)$ such that $(a, \alpha p + (1-\alpha) r) \succ (a,q) \succ (a, \beta p + (1-\beta) r)$.

(A.3) **Conditional Independence -** For each $a \in A$ and all $(a,p), (a,q), (a,r) \in \mathbb{C}$ and $\alpha \in (0,1]$, $(a,p) \succcurlyeq (a,q)$ if and only if $(a, \alpha p + (1-\alpha) r) \succcurlyeq (a, \alpha q + (1-\alpha) r)$.

The next axiom asserts that the decision maker's risk preferences are action-independent.

(A.4) **Action-Independent Risk Preferences -** For all $a, a' \in A$ and $p, q \in \Delta(\mathfrak{C})$, $(a,p) \succcurlyeq (a,q)$ if and only if $(a',p) \succcurlyeq (a',q)$.

The next axiom asserts that the valuations of the actions in A and the distributions in $\Delta(\mathfrak{C})$ are additively separable. To state the axiom I introduce the following additional notations and definitions. Let $\bar{C}, \underline{C} \subset C$ be the subsets of maximal and minimal outcomes in C. Formally, for all $\bar{c} \in \bar{C}$ and $\underline{c} \in \underline{C}$, $(a, \delta_{\bar{c}}) \succ (a,p) \succ (a, \delta_{\underline{c}})$ for all $(a,p) \in A \times \Delta(\mathfrak{C}) \setminus \{\delta_c \mid c \in \bar{C} \cup \underline{C}\}$.[8] If $\bar{C}$ or $\underline{C}$ contain more that one element, choose arbitrarily any one of the maximal and minimal elements.

A pair of actions $a, a' \in A$ is said to be *directly linked* if neither $(a, \delta_{\underline{c}}) \succcurlyeq (a', \delta_{\bar{c}})$ nor $(a', \delta_{\underline{c}}) \succcurlyeq (a, \delta_{\bar{c}})$. They are *indirectly linked* if there exist a sequence of actions $a^1, ..., a^k \in A$ such that $a^1 = a$, $a^k = a'$ and for $i = 1, ..., k-1$, a^i and a^{i+1} are directly linked. For a and a' that are directly linked there is $\alpha' \in (0,1)$ such that $(a', \delta_{\bar{c}}) \sim (a, \alpha' \delta_{\bar{c}} + (1-\alpha') \delta_{\underline{c}})$ or $\alpha \in (0,1)$ such that $(a, \delta_{\bar{c}}) \sim (a', \alpha \delta_{\bar{c}} + (1-\alpha) \delta_{\underline{c}})$. Suppose that a and a' are indirectly linked and let $(a, \delta_{\underline{c}}) \succ (a', \delta_{\bar{c}})$. Then there are $\alpha^i \in (0,1)$ such that $(a^i, \delta_{\bar{c}}) \sim (a^{i+1}, \alpha^i \delta_{\bar{c}} + (1-\alpha^i) \delta_{\underline{c}})$, $i = 1, ..., k-1$.

For each $\alpha \in [0,1]$, let $p_\alpha := \alpha \delta_{\bar{c}} + (1-\alpha) \delta_{\underline{c}}$. Then, for all $a, a' \in A$ that are directly linked there exist $\alpha, \beta \in [0,1]$ such that $(a, p_\alpha) \sim (a', p_\beta)$. That such $\alpha, \beta \in [0,1]$ exist follows from the fact that, since a and a' are directly linked, either $(a', \delta_{\bar{c}}) \succ (a, \delta_{\underline{c}})$ or $(a, \delta_{\bar{c}}) \succ (a', \delta_{\underline{c}})$. Consider the former case (the argument of the latter case is the same). There is an interval $I \subset [0,1]$ such that $(a, \delta_{\bar{c}}) \succ (a', p_\beta) \succ (a, \delta_{\underline{c}})$, for all $\beta \in I$. Then, for every given given $\beta \in I$, by conditional Archimedean, there is $\alpha \in [0,1]$ such that $(a, p_\alpha) \sim (a', p_\beta)$.

[8]That $\bar{c}$ and $\underline{c}$ are independent of a is an implications of (A.4).

The difference $\alpha-\beta$ is a measure of the implicit cost difference between choosing a and a'. The next axiom asserts that this cost difference is independent of the distributions that, together with these actions, constitute the elements of the choice set.

(A.5) **Separability** - For any $a, a' \in A$, that are directly linked, and $\alpha, \beta, \alpha', \beta' \in [0,1]$, $(a, p_\alpha) \sim (a', p_\beta)$ and $(a, p_{\alpha'}) \sim \left(a', p_{\beta'}\right)$ if and only if $\alpha - \beta = \alpha' - \beta'$.

The last axiom asserts that, conditional on the actions, not all the elements of $\Delta(\mathfrak{C})$ and equally preferred.

(A.6) **Non-triviality -** For each $a \in A$, $(a, \delta_{\bar{c}}) \succ \left(a, \delta_{\underline{c}}\right)$.

3.2 Representations

Theorem: *Let $\succcurlyeq$ be a preference relation on $\mathbb{C}$ and suppose that all the alternatives in A are directly or indirectly linked then $\succcurlyeq$ satisfies (A.1) - (A.6) if and only if there exist a real-valued functions u on $\mathfrak{C}$ and ζ on A such that, for all $(a,p), (a',p') \in \mathbb{C}$,*

$$(a,p) \succcurlyeq \left(a', p'\right) \Leftrightarrow \Sigma_{c\in\mathfrak{C}} u(c)\, p(c) + \zeta(a) \geq \Sigma_{c\in\mathfrak{C}} u(c)\, p'(c) + \zeta\left(a'\right). \tag{6}$$

Moreover, the function $u(\cdot) + \zeta(\cdot)$ is unique up to positive linear transformation.

Proof. (a) (Sufficiency) Suppose that $\succcurlyeq$ satisfies (A.1)-(A.4). By the expected utility theorem, $\succcurlyeq$ satisfies (A.1)-(A.3) if and only if there exist real-valued functions $u(\cdot, a)$, $a \in A$, such that, for all $(a,p), (a,p') \in \mathbb{C}$,

$$(a,p) \succcurlyeq \left(a, p'\right) \Leftrightarrow \Sigma_{c\in\mathfrak{C}} u(c,a)\, p(c) \geq \Sigma_{c\in\mathfrak{C}} u(c,a)\, p'(c).$$

Moreover, for each $a \in A$, the function $u(\cdot, a)$ is unique up to positive linear transformation.

Axiom (A.4) implies that, for all $a, a' \in A$, $u(\cdot, a)$ and $u(\cdot, a')$ are positive linear transformations of one another. Fix $\hat{a} \in A$ and let $u(\cdot, \hat{a}) := u(\cdot)$, then for all $a \in A$, $u(\cdot, a) = u(\cdot)\lambda(a) + \zeta(a)$, where, for all $a \in A$, $\lambda(a) > 0$, and $\lambda(\hat{a}) = 1$, $\zeta(\hat{a}) = 0$. Thus, for all $(a,p), (a',p') \in \mathbb{C}$,

$$(a,p) \succcurlyeq \left(a', p'\right) \Leftrightarrow \Sigma_{c\in\mathfrak{C}} u(c)\,\lambda(a)\, p(c) + \zeta(a) \geq \Sigma_{c\in\mathfrak{C}} u(c)\,\lambda\left(a'\right) p'(c) + \zeta\left(a'\right). \tag{7}$$

Let $\bar{c}, \underline{c} \in C$ be maximal elements of C (i.e., $(a, \delta_{\bar{c}}) \succcurlyeq (a,p) \succcurlyeq \left(a, \delta_{\underline{c}}\right)$, for all $(a,p) \in \mathbb{C}$). For each $\alpha \in (0,1)$ define $p_\alpha = \alpha\delta_{\bar{c}} + (1-\alpha)\,\delta_{\underline{c}}$. Suppose that a and a' are directly linked and let $\alpha, \beta, \alpha', \beta' \in [0,1]$ be such that $(a, p_\alpha) \sim (a', p_\beta)$ and $(a, p_{\alpha'}) \sim \left(a', p_{\beta'}\right)$. Then, by (7),

$$[\alpha u(\bar{c}) + (1-\alpha)\, u(\underline{c})]\,\lambda(a) + \zeta(a) = [\beta u(\bar{c}) + (1-\beta)\, u(\underline{c})]\,\lambda\left(a'\right) + \zeta\left(a'\right) \tag{8}$$

and

$$\left[\alpha' u(\bar{c}) + \left(1-\alpha'\right) u(\underline{c})\right] \lambda(a) + \zeta(a) = \left[\beta' u(\bar{c}) + \left(1-\beta'\right) u(\underline{c})\right] \lambda\left(a'\right) + \zeta\left(a'\right). \tag{9}$$

But (8) is equivalent to

$$\left[u(\bar{c}) - u(\underline{c})\right]\left[\alpha\lambda(a) - \beta\lambda\left(a'\right)\right] = u(\underline{c})\left[\lambda\left(a'\right) - \lambda(a)\right] + \zeta\left(a'\right) - \zeta(a), \tag{10}$$

and (9) is equivalent to

$$\left[u(\bar{c}) - u(\underline{c})\right]\left[\alpha'\lambda(a) - \beta'\lambda\left(a'\right)\right] = u(\underline{c})\left[\lambda\left(a'\right) - \lambda(a)\right] + \zeta\left(a'\right) - \zeta(a). \tag{11}$$

Subtracting (9) from (8) we obtain

$$\left[u(\bar{c}) - u(\underline{c})\right]\left[\left(\alpha-\alpha'\right)\lambda(a) - \left(\beta-\beta'\right)\lambda\left(a'\right)\right] = 0. \tag{12}$$

Since a and a' are directly linked, by (A.5), $\alpha - \beta = \alpha' - \beta'$. Hence, $\alpha - \alpha' = \beta - \beta'$. Thus, by (12),

$$\left(\alpha-\alpha'\right)\left[u(\bar{c}) - u(\underline{c})\right]\left[\lambda(a) - \lambda\left(a'\right)\right] = 0. \tag{13}$$

But $(\alpha - \alpha') \neq 0$ and, by (A.6), $u(\bar{c}) - u(\underline{c}) > 0$, implying that $\lambda(a) = \lambda(a')$, for all $a, a' \in A$. In particular, $\lambda(a) = \lambda(\hat{a}) = 1$, for all $a \in A$. Hence, by (7),

$$(a,p) \succcurlyeq \left(a',p'\right) \Leftrightarrow \Sigma_{c\in\mathfrak{C}} u(c)\, p(c) + \zeta(a) \geq \Sigma_{c\in\mathfrak{C}} u(c)\, p'(c) + \zeta\left(a'\right). \tag{14}$$

which is the representation (6).

If a and a' are indirectly linked then $\zeta(a) - \zeta(a') = \Sigma_{i=1}^{k-1}\left(\zeta\left(a^i\right) - \zeta\left(a^{i+1}\right)\right)$, where $\left(a^1,...,a^k\right)$ is a sequence that links $a = a^1$ and $a' = a^k$. Thus, (6) holds with $\zeta(a) = \zeta(a') + \Sigma_{i=1}^{k-1}\left(\zeta\left(a^i\right) - \zeta\left(a^{i+1}\right)\right)$.

(Necessity) The proof that (6) implies that $\succcurlyeq$ satisfies (A.1)-(A.6) is obvious and is, therefore, omitted.

The uniqueness of $u(\cdot) + \zeta(\cdot)$ follows from the uniqueness of $u(\cdot, a)$. ■

The function ζ captures the (utility) cost of the actions and, if the preference relation satisfies (A.1)-(A.6), it is additively separable from the expected utility of the outcomes. Implicit in the representation is the expression $u(\theta)\, p(\theta)$. One should think of θ as representing unanticipated outcomes, or outcomes 'not in C'. Accordingly, $u(\theta)$ captures the decision maker's valuation of discovering outcomes of whose existence he is unaware.

4 Discussion

4.1 Dynamic choice behavior

The potential of discovering new consequences raises novel conceptual issues regarding the modeling of choice dynamics. Specifically, the discovery of unanticipated consequences expands the support of the distributions and, thereby, changes the nature of the risks associated with future actions. Furthermore, unawareness of the consequences that may obtain makes it impossible to design plans of choosing future actions contingent on the realizations of such consequences. Put differently, whereas the probability of encountering an unforeseeable consequences is known, the nature of such consequences, once they become concrete, is inconceivable a priori. This make it impossible to plan a course of action which depends on these consequences, yet, they affect the evaluation of future actions.

To overcome this difficulty, I invoke the fact that, given the structure of the preference relation, once new consequences are discovered, the preference relation extended to the expanded set of known consequences implies that the newly discovered consequences are assigned utilities. And *while the consequences themselves may be inconceivable, the utilities assigned to them are conceivable.* The choices of future actions are contingent on the utilities rather than the consequences themselves.

To grasp the choice dynamics implied by the model, consider a decision problem that requires the choice of actions over two consecutive periods. Suppose that at time $\tau = N$, the action a_i has been taken N_i times, $i = 1, ...m$, the set of observed outcomes is C_τ and the corresponding vector of frequency distribution is $\mathbf{p} = (\mathbf{p}_1, ..., \mathbf{p}_m)$. A choice of $a_i \in A$ induces a conditional probability distribution $\mathbf{p}_\tau(\cdot \mid a_i) \in \Delta(\mathfrak{C}_\tau)$, given by

$$\mathbf{p}_\tau(c \mid a_i) = \Pr\{Z_{i(N_i+1)} = c \mid \mathbf{p}\},$$

for all $c \in C_\tau$, and

$$\mathbf{p}_\tau(\theta_i \mid a_i) = \Pr\{Z_{i(N_i+1)} = \theta_i \mid \mathbf{p}\},$$

where $\Pr\{Z_{i(N_i+1)} = c \mid \mathbf{p}\}$ and $\Pr\{Z_{i(N_i+1)} = \theta_i \mid \mathbf{p}\}$ are given by (4) and (5), respectively.

Consider a choice of a_i, and supposed that the consequence that obtains is $Z_{i(N_i+1)} = c_k$.

(a) If $c_k \in C_\tau$ then $C_{\tau+1} = C_\tau$ and $\mathfrak{C}_{\tau+1} = \mathfrak{C}_\tau$. The decision maker updates the frequency distribution $\mathbf{n}_i(C_\tau)$ to $\mathbf{n}'_i(C_{\tau+1}) = \left(n_{i1}, ..., n_{ik} + 1, ..., n_{i|C|}\right)$ and, for all $j \neq i$, $\mathbf{n}_j(C_\tau) = \mathbf{n}'_j(C_{\tau+1})$. The corresponding probability vector, $\mathbf{p}_\tau(a_i, c_k) = (\mathbf{p}'_1, .., \mathbf{p}'_k, ..., \mathbf{p}'_m)$, is given by:

$$\mathbf{p}'_i = (p'_{ic_1}, ..., p'_{ic_{|C|}}, p'_{i\theta_i}) = \left(n_{i1}, ..., n_{ik} + 1, ..., n_{ic_{|C|}}, p_{i\theta_i}\right) / \left(N_i + 1 + \theta_i\right), \tag{15}$$

and $\mathbf{p}'_j = \mathbf{p}_j,\ j \neq i$.

Define

$$\kappa'\left(i,j\right) = \frac{\langle \mathbf{p}'_i, \mathbf{p}'_j \rangle}{\parallel \mathbf{p}'_i \parallel \parallel \mathbf{p}'_j \parallel},\ i,j \in \{1, ..., m\}.$$

Letting $N'_i = N_i + 1$ and, for $j \neq i$, $N'_j = N_j$. Then, for all $c \in \mathfrak{C}_{\tau+1}$ and $j = 1, ..., m$,

$$\Pr\{Z_{i\left(N'_i+1\right)} = c \mid \mathbf{p}_{\tau,} \left(a_i, c_k\right)\} = \sum_{j=1, j\neq i}^{m} p'_{jc}\kappa'\left(i,j\right)\frac{N'_j}{N+1} + p'_{ic}, \tag{16}$$

and

$$\Pr\{Z_{j\left(N'_j+1\right)} = c \mid \mathbf{p}_{\tau}, \left(a_i, c_k\right)\} = \sum_{r=1, r\notin\{i,j\}}^{m} p'_{rc}\kappa'\left(j,r\right)\frac{N_r}{N+1} + p'_{ic}\kappa'\left(j,i\right)\frac{N_i+1}{N+1} + p'_{jc}. \tag{17}$$

Given that a_i was chosen in time τ resulting in the outcome $c_k \in C_{\tau+1}$, a choice of a_j at time $\tau + 1$ induces a conditional probability distribution $\mathbf{p}_{\tau+1}\left(\cdot \mid a_j\right) \in \Delta\left(\mathfrak{C}_{\tau+1}\right)$, given by

$$\mathbf{p}_{\tau+1}\left(c \mid a_j\right) = \Pr\{Z_{j\left(N'_j+1\right)} = c \mid \mathbf{p}_{\tau}, \left(a_i, c_k\right)\},$$

for all $c \in \mathfrak{C}_{\tau+1}$.

Given the choice of a_τ in the first period, in the second period the decision maker chooses an action contingent on the consequence c'. Formally,

$$a^*\left(a_\tau, c'\right) \in \arg\max_A \left[\Sigma_{c\in\mathfrak{C}_{\tau+1}} u\left(c\right)\mathbf{p}_{\tau+1}\left(c \mid a, \left(a_\tau, c'\right)\right) + \zeta\left(a\right)\right].$$

(b) If $Z_{i(N_i+1)} = \theta_i$, then the set of known consequences is augmented by the addition of a newly discovered outcome, $\hat{c} \notin C_\tau$. Formally, $C_{\tau+1} = C_\tau \cup \{\hat{c}\}$ and $\mathfrak{C}_\tau \subset \mathfrak{C}_{\tau+1} = C_{\tau+1} \cup \{\theta\}$. Because $\hat{c}$ is unforeseeable, it is impossible to decide ahead of time on a plan of actions contingent on its realizations. However, a decision maker whose preference relation on the augmented choice set $\mathbb{C}_{\tau+\nVdash} = A \times \mathfrak{C}_{\tau+1}$, is depicted by (A.1)-(A.6), anticipates being able to assign utility value to the unforeseen consequences, whatever it may happened to be. Put differently, the decision maker anticipates assigning a real number, $u(\hat{c})$ to every $\hat{c}$ that may obtain. The crucial point is, *the decision maker cares about the value of the* *utility rather than the particular consequence that yields it.* Consequently, the decision maker foresees choosing, in the second period, the action $a^* (a_\tau, u)$ given by

$$a^* (a_\tau, u) \in \arg\max_A [\Sigma_{c \in \mathfrak{C}_\tau} u (c) \mathbf{p}_{\tau+1} (c \mid a, a_\tau) + u\mathbf{p}(u \mid a, a_\tau) + \zeta (a)] ,$$

where u is the utility of the unforeseeable consequence when it takes a concrete shape.

Viewed form the first period (i.e., before the new consequence are discovered) the probability $\mathbf{p}(u \mid a, a_\tau)$ may be decomposed as follows:

$$\mathbf{p}(u \mid a, a_\tau) = \Pr(\theta \mid a_\tau) \times \Pr(\theta = u \mid a, a_\tau)$$

where $\Pr(\theta \mid a_\tau) = \Sigma_c \in \mathfrak{C}_\tau \mathbf{p}(c \mid a_\tau)$ and the *utility risk*, $\Pr(\theta = \cdot \mid a, a_\tau)$ (i.e., the probability that the unforeseen consequence yield a particular utility value) is *subjectively assessed* by the decision maker. In particular, if the decision maker is guided by his pervious experience he may assign to each utility value the probability that that value obtained in the past. Formally,

$$\Pr(\theta = u \mid a, a_\tau) = \Sigma_{c \in \{c' \in C_\tau \mid u(c') = u\}} p(c \mid a, a_\tau).$$

Define

$$V(a^*(a_\tau, u)) = \Sigma_{c \in \mathfrak{C}_\tau} u (c) \mathbf{p}_{\tau+1} (c \mid a^* (a_\tau, u) , a_\tau) + u\mathbf{p}(u \mid a, a_\tau) + \zeta (a^* (a_\tau, u)) ,$$

Then the first-period decision problem is: Choose $a_\tau \in A$ so as to maximize

$$\Sigma_{c' \in \mathfrak{C}_\tau} \left[u \left(c'\right) + \lambda[\Sigma_{c \in \mathfrak{C}\tau}[u (c) + \zeta \left(a^* \left(a_\tau, c'\right)\right)]\mathbf{p}_{\tau+1} \left(c \mid a^* \left(a_\tau, c'\right)\right) \mathbf{p}(c' \mid a_\tau) + \right.$$

$$\Pr(\theta \mid a_\tau) \int_{-\infty}^{\infty} (V(a^*(a_\tau, u)) + \zeta (a^*(a_\tau, u) du) \Pr(\theta = u \mid a, a_\tau) + \zeta (a_\tau) ,$$

where $\lambda \in [0, 1]$ denotes the discount rate.

The choice of the first-period action yields payoffs, in the form of outcomes, and information regarding the probabilistic payoffs of actions, including potential discovery of unanticipated outcomes. This dual role implies that the first period choice involves exploitation-exploration trade-off. In other word, it may be that $(a, \mathbf{p}_\tau (\cdot \mid a_\tau)) \not\succcurlyeq (a', \mathbf{p}_\tau (\cdot \mid a'_\tau))$ and yet, a'_τ is chosen in the first period if it is more informative about the distribution of outcomes in the second period.

4.2 Multi-armed bandit

A strand of literature displaying some of the ingredients of the model of this paper deals with multi-armed bandit problems. In its most familiar form it is a sequential decision problem that requires the decision maker (e.g., gambler) choose a sequence (finite or infinite) of arms, of distinct slot machines, to pull so as to maximize the expected present value of his reward. The distributions of the payoffs of the different arms are unknown. Each choice of arm pays off immediately and, at the same time, informs the player about the distribution of the payoffs associated with the arm. The most common variation of the multi-arm bandit problems assumes that the random returns of the distinct arms are stochastically independent. Other variations include correlated random payoffs across arms. In either case, since the possible payoffs are supposed to be known, the learning takes the form of updating the distributions by the application of Bayes' rule.[9]

The main differences between the multi-arm bandit models and the model of this paper are: (a) Whereas in the multi-armed bandit problem it is assumed that the set of possible payoffs is known and fixed, the focal issue of this paper is the process of discovery of unanticipated outcomes, or payoffs, and (b) A consequence of (a) is that unlike the exploration in the multi-armed bandit game, which consists of updating the distributions of the arms by the application of Bayes' rule, the exploration in the model of this paper includes both the discovery of new, unanticipated, outcomes and the updating of the probability distributions of the known outcomes. This former aspect renders Bayes' rule inapplicable. Instead, learning is accomplished by the application of Ewens' (1972) generalization of De Morgan's rule.

4.3 Related literature

Schipper (2022) derives the predictive probabilities of the De Morgan rule, Ewens sampling rule, as subjective probabilities. In particular, Schipper considers the process of repeated sampling from a population, using the same sampling procedure, and studies the question of what must be true about the pattern of a decision maker's betting on outcomes (including the discovery of novelty) for it to display beliefs that agree with these rules. Whereas the main concern of this paper is the modeling of the behavior of decision makers whose beliefs are represented by Ewens' (1972) sampling rule, the main thrust of Schipper's

[9] Bergemann and Välimäki (2008) provide a breif review of this literature.

work, the characterization the subjective beliefs that agree with the objective predictions of the exchangeable random partition models and display ‘reverse Bayesianism’ á la Karni and Vierø (2013, 2017). In particular, Schipper shows that models exhibiting ‘reverse Bayesianism’ include, among others, the De Morgan model and some variations of it, including Ewens’ sampling rule.

The exploitation-exploration aspect of the dynamic application of the model is a feature shared by Karni (2022). Unlike the present paper, in which the predictions of probabilities of the action-contingent outcomes is arrived at by induction, the probabilities of the outcomes in Karni (2022) are predicted by theoretical models. Moreover, whereas in this paper actions may discover unanticipated outcomes, in Karni (2022) the set of outcomes is known and fixed, so the exploration aspect is captured by the updating of the decision maker’s probabilistic belief in the validity of the theories using Bayes’ rule.

Unlike the Bayesian approach to unawareness modeled in Karni and Vierø (2013, 2017), in which the probability assigned to newly discovered events in the state space and the residual probability assigned to unanticipated outcomes, are subjective and derived from the preference relation. The probability assigned to unanticipated outcomes in the model of this paper is objective and follows algorithmic process. Moreover, the approach of Karni and Vierø allows the probability of unanticipated outcomes to increase with the discovery of new consequences. This possibility is not allowed by the model of this paper. The probability of unforeseen consequences must decay as the number of observations increases.

Invoking the state-space formulation of increasing awareness introduced by Karni and Vierø (2013, 2017), Grant et al. (2022) propose a model of learning in which decision makers become aware of new states (i.e., resolutions of uncertainty) through the discovery of unknown actions and consequences. Whereas the observations in this work are action-outcome pairs, the sampling process in the model of Grant et al. yields observations of states restricted by the known acts and consequences at each point in time. Moreover, unlike this paper, in which the underlying stochastic process generating the observations is not assumed to have a specific structure, Grant et al. (2022) assume that the data is generated by a Dirichlet process, which govern the evolution of the decision maker’s beliefs.[10] Finally, the main thrust of their work is the characterization of the learning process as opposed to that of the decision making process.

[10] A focal issue in their work is the ambiguity surounding the base measure of the Dirichlet process

Eichberger and Gouerdjikova (2024) propose a model of decision making under uncertainty whose primitives are cases, (i.e., triplet of action taken, outcome obtained and characteristics, describing background data that determines the outcome corresponding to each and every action). The set of cases is the data available at the point at which a decision maker must choose the next action. The data permits the construction of the relative frequencies of the already known outcomes and characteristics. The set of characteristics in the data may not be complete. It is completed by including a "place holder" signifying not yet-observed characteristic. They axiomatize a representation of preferences that includes a subjective weight assigned to this "place holder", which they interpret as measure of the awareness of unawareness of possible characteristics. In the model of this paper, the set of observed act-outcome pairs is the data and the states are resolutions of the uncertainty, has the flavor of characteristics set in Eichberger and Gouerdjikova (2024). Despite the similar interpretation, however, there is a fundamental distinction between the state-space and the set of characteristics. First, the state space is a derived concept whereas the characteristics are taken to be directly-observed primitive in Eichberger and Gouerdjikova's model. Second, the evolution of the state space is endogenous, following naturally the process of taking actions and discovering new, unanticipated, outcomes, whereas the discovery of new characteristics is exogenous and independent of the actions taken. Third, in the present model the probability of discovering an outcome, not yet seen, (i.e., the measure of unawareness) is monotonic decreasing as a function of the number of actions taken (i.e., that sample size), whereas in the model of Eichberger and Gouerdjikova a discovery of not yet-seen characteristic increases their measure of the unawareness. Forth, unlike the extended set of characteristics whose completion entails a "place holder," the specification of state space does not require such device. The unawareness in this model is about the possible outcomes in which the mutator plays the role of "place holder." Fifth, in the model of this paper the probability distributions on the evolving state space are objectively derived from the data, while the distribution on the set of extended characteristics incorporates a subjective "degrees of unawareness", which is an ingredient of the representation of the preference relations.

References

[1] Arrow, Kenneth J (1965) *Essays on the Theory of Risk Bearing* North Holland, Amsterdam.

[2] Becker, C. K., Melkonyan, T., Proto, E., Soanos, A., Trautman, S. T. (2020) "Reverse Bayesianism: Revising Beliefs in Light of Unforeseen Events," University of Heidelberg.

[3] Bergemann, D. and Välimäki, J. (2008). Bandit problems. In: Durlauf, Steven N., Blume, Lawrence E. (Eds.), 2nd ed.. The New Palgrave Dictionary of Economics, vol. 1. Macmillan Press, pp. 336–340.

[4] Chakravarty, S., Kelsey, D., Teitelbaum, J. C. (2022). "Reverse Bayesianism and Act Independence," *Journal of Economic Theory,* 203, 105495.

[5] De Morgan, A. (1838) *An essay on probabilities, and their application to life contingencies and insurances offices,* London: Longman, Orne, Brown, Green, and Longmans.

[6] Dominiak, A., Tserenjigmid, G. (2018). "Belief Consistency and Invariant Risk Preferences," *Journal of Mathematical Economics* 79, 157-162.

[7] Eichberger, J. and Guerdjikova, A. (2024) "Data, Cases and States," unpublished manuscript.

[8] Ewens, W. J. (1972) "The Sampling Theory of Selectively Neutral Alleles," *Theoretical Population Biology* 3, 87–112.

[9] Grant, S., Quiggin, J. (2013) "Inductive Reasoning about Unawareness," *Economic Theory* 54, 717-755.

[10] Grant, S., Meneghel, I., Tourky, R. (2022) "Learning Under Unawareness," *Economic Theory*

[11] Heifetz, A., Meier, M., Schipper, B. C. (2006) "Interactive Unawareness," *Journal of Economic Theory* 130, 78-94.

[12] Heifetz, A., Meier, M., Schipper, B. C. (2008) "A Canonical Model for Interactive Unawareness," *Games and Economic Behavior* 62, 305-324.

[13] Heifetz, A., Meier, M., Schipper, B. C. (2013) "Unawareness, Beliefs, and Speculative Trade," *Games and Economic Behavior* 77, 100-121.

[14] Karni, E. (2022) "A Theory-Based Decision Model," *Journal of Economic Theory,* 201,

[15] Karni, E., and Schmeidler, D. (1991) "Utility Theory with Uncertainty," in Werner Hildenbrand and Hugo Sonnenschein, eds., *Handbook of Mathematical Economics* vol. IV. Elsevier Science Publishers B.V.

[16] Karni, E., Vierø, M.-L. (2013). "'Reverse Bayesianism': A choice-based theory of growing awareness," *American Economic Review* 103, 2790-2810.

[17] Karni, E., Vierø, M.-L. (2015). "Probabilistic Sophistication and Reverse Bayesianism," *Journal of Risk and Uncertainty*, 50 (2015) 189-208.

[18] Karni, E., Vierø, M.-L. (2017). "Awareness of Unawareness: A theory of Decision Making in the Face of Ignorance," *Journal of Economic Theory* 168, 301-328.

[19] Karni, E., Valenzuela-Stookey, Q., Vierø, M.-L. (2021). "'Reverse Bayesianism': A Generalization," *B. E. Journal of Theoretical Economics* 21, 557-569.

[20] Schipper, B. C. (2022) "Predicting the Unpredictable under Subjective Expected Utility" unpublished manuscript.

[21] Schmeidler, D. and Wakker, P. (1987) "Expected Utility and Mathematical Expectation," in John Eatwell, Murray Milgate, and Peter Newman, eds., *The New Palgrave: A Dictionary of Economics.* Macmillan Press.

[22] Vierø, M.-L. (2021) "An Intertemporal Model of Growing Awareness," *Journal of EconomicTheory* 197, 105351.

[23] Zabell, S. L. (1992) "Predicting the Unpredictable," *Synthese* 90, 205-232.