

Decisions and Discovery: A non-Bayesian Perspective

Edi Karni ^{*}

Johns Hopkins University

July 8, 2024

Abstract

The paper explores a new, non-Bayesian, approach to modeling decision makers' awareness of unawareness of the potential outcomes of their actions. The new approach is an adaptation of Ewens (1972) generalization of De Morgan's (1838) formula of the probability of known and unknown outcomes to the context of decision making under uncertainty, and incorporating it into 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

Decision making is the exercise of choice 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 become aware of a consequence only after it obtains for the first time (i.e., it is unknown unknown). Repeated discovery of unanticipated, novel, consequences, however, alerts the decision maker to the possibility that there may exist additional, unanticipated, consequences whose nature he is unable to conceive of. In other words, the decision maker may be aware of his unawareness, and this awareness of unawareness may affect his choice behavior. Based on their notion of 'reverse Bayesianism',¹ Karni and Vierø (2017) proposed a model of decision making under uncertainty in which decision makers' awareness of their unawareness impact their choice behavior.

In this paper I explore a different, non-Bayesian, approach to modeling decision making under uncertainty, based on inductive inference. The proposed model accommodates decision makers beliefs regarding potential consequences whose existence is not conceivable of before they obtain. To motivate this endeavour, I consider instances requiring making decision under uncertainty, in which decision makers have access to data that may be used to calibrate the likelihoods of the outcomes of their actions. For example, 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 their beliefs. 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 make use the accumulated evidence to form beliefs about the likelihoods of occurrence of known outcomes and potential existence of unforeseeable health effects. Similarly, statins have been used to reduce cholesterol levels for decades. Their side effects, including headache, dizziness, feeling unusually tired or physically weak, digestive system problems, such as constipation, diarrhoea, or indigestion, muscle pain, sleep problems, low

¹See Karni and Vierø (2013).

blood platelet count, are well known, leaving little chance to discovering new, unanticipated, side effects, both in the short and the long run. Clinical trials have shown that a new medication is more effective in lowering the cholesterol levels and has fewer side effects. However, the lack of long-term 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.

This example illustrates the need for theories of decision making under uncertainty, founded on inductive inference, that accommodates the potential existence of unknown unknown 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 the consequences that may result from their choice of actions, and may also be aware of this unawareness. 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 (2017), Dominiak and Tserenjigmid (2018), Karni, Vierø, Valenzuela-Stookey (2021), Chakravarty, Kelsey, and Teitelbaum (2021), Vierø (2021)².

A different, 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 an event obtains. In other words,

“How can we predict the occurrence of something we neither know, nor even suspect, exists? Subjective probability and Bayesian inference, despite their many impressive successes, would seem at a loss to handle such a problem

²The 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).

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, \dots, Z_N of N trials (i.e., observations) t categories, or outcomes, labeled c_1, \dots, c_t , have been observed, then the probability of seeing the outcome on trial $N + 1$ fall into the j -th category is:

$$\Pr\{Z_{(N+1)} = c_j \mid \mathbf{n}\} = \frac{n_j + 1}{N + t + 1}, \quad (1)$$

where $\mathbf{n} = (n_1, \dots, n_t)$ denote the number of times each of the t outcomes occurred in N trials. Notice that $\sum_{j=1}^t \Pr\{Z_{(N+1)} = c_j \mid \mathbf{n}\} < 1$ (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

$$\Pr\{Z_{(N+1)} = \hat{c} \mid \mathbf{n}\} = \frac{1}{N + t + 1}. \quad (2)$$

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

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

⁴See Zabell (1992).

Let n_j be the number of times a ball of color c_j is drawn in N trials and let $\mathbf{n} = (n_1, \dots, n_t)$ be the frequency distribution of the known colors. Then, under exchangeability, \mathbf{n} is a sufficient statistics for the $N = \sum_{j=1}^t n_j$ observations. In other words, under exchangeability all the sequences of observations whose frequency distribution is \mathbf{n} are equally probable. 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

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

In this paper I explore the application of the Ewens’ process to situations in which decision making in the face of uncertainty allows for the possible existence of unanticipated consequences or outcomes. This objective requires the modification and extension of De Morgan’s proposal. In particular, 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 changing his awareness of the possible outcomes of all courses of action. The correlation among the samples from distinct urns necessitates 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 the 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 Morgan proposal. Section 3 introduces a decision making model. Section 4 discusses the implications of a decision model that incorporates the extended De Morgan process. Section 5 includes additional remarks and a brief discussion of related literature.

⁵At 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).

2 Generalized Sampling Process

The extended sampling process is best described by the following multi-urns model that allows for the possibility that the mutator assumes different weights (e.g., different numbers of mutators in the different urns) and the random processes of draws from distinct urns may not be stochastically independent.

Let $\mathcal{U} := \{U_1, \dots, U_m\}$ be a finite set of *urns*. Let N_i be the number of draws from U_i , $i = 1, \dots, m$, containing balls of t_i of different colors and some black balls called mutators. Let θ_i denote the weight of the mutator in urn U_i . Consider the following process. Select an urn from \mathcal{U} and drawn a ball at random from the selected urn. If a colored ball is drawn, then it is replaced together with another ball of the same color. If a mutator is drawn, then it is replaced together with another ball of new color.

Let n_{ik} be the number of balls of color c_k in N_i draws from U_i . Define $C_i = \{c_{i1}, \dots, c_{it_i}\}$, $i = 1, \dots, m$, 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 $N = \sum_{i=1}^m N_i$ draws from all the urns. Let $n_{i\ell}$ denote the number of balls of color $c_\ell \in C$ in N_i draws from U_i . Denote by $\mathbf{n}_i(C) = (n_{i1}, \dots, n_{i|C|})$ the frequency of draws of the known colors after a sequence $Z_{i1}, Z_{i2}, \dots, Z_{iN_i}$ of draws from U_i . Note that $n_{i\ell} = 0$ if $c_\ell \in C \setminus C_i$.

Given N_1, \dots, N_m and C , let $p_{ic_k} := n_{ic_k} / (N_i + \theta_i)$, $k \in \{1, \dots, |C|\}$, and $p_{i\theta_i} := \theta_i / (N_i + \theta_i)$. Clearly, $p_{ic_k} = 0$ if $c_k \in C \setminus C_i$. Let $\mathbf{p}_i := (p_{ic_1}, \dots, p_{ic_{|C|}}, p_{i\theta_i})$, $i = 1, \dots, m$. For every pair of probability vectors $\mathbf{p}_i, \mathbf{p}_j$, $i, j \in \{1, \dots, m\}$ 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, $\kappa(i, j) = \cos \tau$ is a measure the angle between the two probability vectors. I interpret it to be a measure of *similarity* between the underlying stochastic processes. 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 $i = 1, \dots, m$, (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 $\lim_{N_i \rightarrow \infty} \kappa(i, j) = 0$, for all $i, j = 1, \dots, m$. 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 $i, j = 1, \dots, m$), in

the sense that nothing is known about their content.

It is worth mentioning that a mathematical property of the similarity measure is that it 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 $\hat{\mathbf{p}}_j := \mathbf{p}_j / (1 - p_{j\theta_j})$ are the same. However, the larger is the difference in the sample sized (i.e., the weight 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 , $i = 1, \dots, m$, the probability of observing $c \in C$ conditional on a draw from U_i is:

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

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

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

In the analysis that follows I assume the decision maker predicts the outcomes of his actions using these formulas. If $c \in C \setminus C_i$ then it is known to exist 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

⁶De Morgan's and Ewen's prediction rules are characterized by partition exchangeability. The prediction 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 exchangeability, whereas the predictive rule corresponding to flipping distinct, correlated, coins do not satisfy exchangeability.

3 The Decision Model

3.1 The choice set and the structure of the preference relations

Let $A := \{a_1, \dots, a_m\}$ 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 $U_i \in \mathcal{U}$. 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 and may be inconceivable. Let $\Delta(\mathfrak{C})$ denote the set of probability distributions on \mathfrak{C} , where $\Pr\{\theta\} = 1 - \sum_{c \in C} \Pr\{c\}$.

The product set $\mathbb{C} = A \times \Delta(\mathfrak{C})$ is said to be the *choice set*. A binary relation \succsim on \mathbb{C} , is a *preference relation*. Denote by \succ and \sim the asymmetric and symmetric parts of \succsim , 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 predictive probabilities in $\Delta(\mathfrak{C})$ are obtained from the formulas (4) and (5). The structure of the preference relation is depicted axiomatically as follows.

(A.1) **Weak Order (WO)** - \succsim is complete and transitive.

(A.2) **Conditional Archimedean (CA)** - 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 (CI)** - For each $a \in A$ and all $(a, p), (a, q), (a, r) \in \mathbb{C}$ and $\alpha \in (0, 1]$, $(a, p) \succsim (a, q)$ if and only if $(a, \alpha p + (1 - \alpha)r) \succsim (a, \alpha q + (1 - \alpha)r)$.

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

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

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 $(a, p) \in A \times \Delta(\mathfrak{C}) \setminus \{\delta_c \mid c \in \bar{C} \cup \underline{C}\}$.⁷ If \bar{C} or \underline{C} contain more than 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 $(a, \delta_{\underline{c}}) \succsim (a', \delta_{\bar{c}})$ nor

⁷That \bar{c} and \underline{c} are independent of a is an implications of (A.4).

$(a', \delta_{\underline{c}}) \succcurlyeq (a, \delta_{\bar{c}})$. They are *indirectly linked* if there exist a sequence of actions $a^1, \dots, a^k \in A$ such that $a^1 = a$, $a^k = a'$ and for $i = 1, \dots, k-1$, 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 $(a^i, \delta_{\bar{c}}) \sim (a^{i+1}, \alpha^i \delta_{\bar{c}} + (1 - \alpha^i) \delta_{\underline{c}})$, $i = 1, \dots, k-1$.

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 $\beta \in I$, by conditional Archimedean, there is $\alpha \in [0, 1]$ such that $(a, p_\alpha) \sim (a', p_\beta)$.

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 $(a, p_{\alpha'}) \sim (a', p_{\beta'})$ 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 $a \in A$, $(a, \delta_{\bar{c}}) \succ (a, \delta_{\underline{c}})$.

3.2 Representations

Theorem: Let \succcurlyeq be a preference relation on \mathfrak{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 $(a, p), (a', p') \in \mathfrak{C}$,

$$(a, p) \succcurlyeq (a', p') \Leftrightarrow \sum_{c \in \mathfrak{C}} u(c) p(c) + \zeta(a) \geq \sum_{c \in \mathfrak{C}} u(c) p'(c) + \zeta(a'). \quad (6)$$

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 $(a, p), (a, p') \in \mathbb{C}$,

$$(a, p) \succ (a, p') \Leftrightarrow \sum_{c \in \mathfrak{C}} u(c, a) p(c) \geq \sum_{c \in \mathfrak{C}} u(c, a) p'(c).$$

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 $(a, p), (a', p') \in \mathbb{C}$,

$$(a, p) \succ (a', p') \Leftrightarrow \sum_{c \in \mathfrak{C}} u(c) \lambda(a) p(c) + \zeta(a) \geq \sum_{c \in \mathfrak{C}} u(c) \lambda(a') p'(c) + \zeta(a'). \quad (7)$$

Let $\bar{c}, \underline{c} \in C$ be maximal elements of C (i.e., $(a, \delta_{\bar{c}}) \succ (a, p) \succ (a, \delta_{\underline{c}})$, 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 $(a, p_{\alpha'}) \sim (a', p_{\beta'})$. Then, by (7),

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

and

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

But (8) is equivalent to

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

and (9) is equivalent to

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

Subtracting (9) from (8) we obtain

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

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

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

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),

$$(a, p) \succcurlyeq (a', p') \Leftrightarrow \sum_{c \in \mathfrak{C}} u(c) p(c) + \zeta(a) \geq \sum_{c \in \mathfrak{C}} u(c) p'(c) + \zeta(a'). \quad (14)$$

which is the representation (6).

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

(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 Behavioral Implications

As the examples in the introduction suggest, the approach taken in this paper may be interpreted as a model depicting one choice conditional on the data currently available about the outcomes corresponding to actions taken in the past (e.g., the vaccine effects), or a model of repeated choice based on the cumulative experience of selecting alternative actions (e.g., the experience of dining in different restaurants).

4.1 Static Choice Behavior

Consider first the case of a single choice. Given N_1, \dots, N_m , and a set C of observed outcomes, denote the corresponding vector of frequency distribution by $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_m)$. Choosing $a \in A$ induces a conditional probability distribution $\mathbf{p}(\cdot | a) \in \Delta(\mathfrak{C})$, given by

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

for all $c \in C$, and

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

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.

By the theorem, if a preference relation on \mathbb{C} satisfies the axioms (A.1)-(A.6) then the decision problem is to choose $a \in A$ so as to maximize $\sum_{c \in \mathbb{C}} u(c) \mathbf{p}(c \mid a) + \zeta(a)$. This representation includes the decision maker's assessment of the probability, $\mathbf{p}(\theta_i \mid a_i)$, of encountering outcomes of whose existence he is unaware and the utility evaluation of such occurrences.

4.2 Dynamic Choice Behavior

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

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

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

and $\mathbf{p}'_j = \mathbf{p}_j$, $j \neq i$. If the outcome is $Z_{i(N_i+1)} = \theta$, then the set of known outcomes is augmented by the addition of the newly discovered outcome, $\hat{c} \notin C_\tau$. Formally, $C_{\tau+1} = C_\tau \cup \{\hat{c}\}$.

Define

$$\kappa'(i, j) = \frac{\langle \mathbf{p}'_i, \mathbf{p}'_j \rangle}{\|\mathbf{p}'_i\| \|\mathbf{p}'_j\|}, \quad i, j \in \{1, \dots, m\}.$$

Letting $N'_i = N_i + 1$ and, for $j \neq i$, $N'_j = N_j$. Then, for all $c \in C_{\tau+1} \cup \{\theta\}$, and $j = 1, \dots, m$,

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

and

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

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 $\mathbf{p}_{\tau+1}(\cdot \mid a_j) \in \Delta(\mathfrak{C})$, given by

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

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

Denote by a_τ and c' the action chosen at time τ the outcome that results. By the theorem, if a preference relation on \mathbb{C} satisfies the axioms (A.1)-(A.6) then, for all $a, a' \in A$,

$$\begin{aligned} (a, \mathbf{p}_{\tau+1}(\cdot \mid a, a_\tau, c')) &\succsim (a', \mathbf{p}_{\tau+1}(\cdot \mid a', a_\tau, c')) \Leftrightarrow & (18) \\ \sum_{c \in \mathfrak{C}_{\tau+1}} u(c) \mathbf{p}_{\tau+1}(c \mid a, a_\tau, c') + \zeta(a) &\geq \sum_{c \in \mathfrak{C}_{\tau+1}} u(c) \mathbf{p}_{\tau+1}(c \mid a', a_\tau, c') + \zeta(a'). \end{aligned}$$

Consequently, in the second and final period the decision maker chooses the action

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

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

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

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

The choice of the first-period action yields a payoff in the form of an outcome and, simultaneously, and information regarding the probabilistic payoffs of actions including the potential discovery of unanticipated outcomes. This dual role implies that first period choice involves exploitation-exploration trade-off. In other word, it may be that $(a, \mathbf{p}_\tau(\cdot \mid a_\tau)) \succsim (a', \mathbf{p}_\tau(\cdot \mid a'_\tau))$ and yet, a'_τ is chosen in the first period if it is more informative about the distribution of outcomes in the second period, that is, if

$$\begin{aligned} &\arg \max_A \sum_{c' \in \mathfrak{C}_\tau} [\sum_{c \in \mathfrak{C}_{\tau+1}} u(c) \mathbf{p}_{\tau+1}(c \mid a^*(a'_\tau, c')) + \zeta(a^*(c'))] \mathbf{p}_\tau(c' \mid a'_\tau) \\ &> \arg \max_A \sum_{c' \in \mathfrak{C}_\tau} [\sum_{c \in \mathfrak{C}_{\tau+1}} u(c) \mathbf{p}_{\tau+1}(c \mid a^*(a_\tau, c')) + \zeta(a^*(c'))] \mathbf{p}_\tau(c' \mid a_\tau). \end{aligned}$$

5 Concluding Remarks

5.1 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 an element, a *true state*, that once known, the consequences of every conceivable course of action are known.

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 C denote the set of known outcomes, or consequences, after N choices of actions in A . The state space describing the uncertainty before the $N + 1$ choice consists of all the mappings from the set of actions to the set of known and unknown outcomes (i.e., $S = \{s : A \rightarrow \mathfrak{C}\}$).⁸ Clearly, this definition of states represents the resolution, in the sense defined above, of the uncertainty the decision maker faces before choosing his next action. Moreover, given N_1, \dots, N_m , $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_m)$ and C , the probability of the state $s = (s(a_1), \dots, s(a_m))$ is:

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

Obviously, the state space in this model is neither primitive nor immutable. In fact, once a new outcome 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: $S' = \{s' : A \rightarrow \mathfrak{C}'\}$. The probabilities of the states are:

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

where $\hat{\mathbf{p}} = (\hat{\mathbf{p}}_1, \dots, \hat{\mathbf{p}}_m)$ and $\hat{\mathbf{p}}_i = (\hat{p}_{i1}, \dots, \hat{p}_{i|C'|}, \hat{p}_{i\theta_i}), i = 1, \dots, m$, 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,

⁸This notion of states was describe in Schmeidler and Wakker (1987) and Karni and Schmeidler (1991) and was invoked in Karni and Viero (2013), (2017).

let $\mathbf{p}' = (\mathbf{p}'_1, \dots, \mathbf{p}'_m)$, where $\mathbf{p}'_i, i = 1, \dots, m$, are given by (15) then

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

5.2 Related literature

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) 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.

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.

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 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.

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.⁹ Finally, the main thrust of their work is the characterization of the learning process as opposed to that of the decision making 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 in which a decision maker must choose the next action. The data permits the construction of the relative frequencies of the already 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

⁹A focal issue in their work is the ambiguity surrounding the base measure of the Dirichlet process

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 the present paper, the set of observed act-outcome pairs may be the data and the state-space, whose elements 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. Fourth, unlike the set 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] 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.
- [2] Chakravarty, S., Kelsey, D., Teitelbaum, J. C. (2022). “Reverse Bayesianism and Act Independence,” *Journal of Economic Theory*, 203, 105495.
- [3] De Morgan, A. (1838) *An essay on probabilities, and their application to life contingencies and insurances offices*, London: Longman, Orne, Brown, Green, and Longmans.
- [4] Dominiak, A., Tserenjigmid, G. (2018). “Belief Consistency and Invariant Risk Preferences,” *Journal of Mathematical Economics* 79, 157-162.
- [5] Eichberger, J. and Guerdjikova, A. (2024) “Data, Cases and States,” unpublished manuscript.
- [6] Ewens, W. J. (1972) “The Sampling Theory of Selectively Neutral Alleles,” *Theoretical Population Biology* 3, 87–112.
- [7] Grant, S., Quiggin, J. (2013) “Inductive Reasoning about Unawareness,” *Economic Theory* 54, 717-755.
- [8] Grant, S., Meneghel, I., Tourky, R. (2022) “Learning Under Unawareness,” *Economic Theory*
- [9] Heifetz, A., Meier, M., Schipper, B. C. (2006) “Interactive Unawareness,” *Journal of Economic Theory* 130, 78-94.
- [10] Heifetz, A., Meier, M., Schipper, B. C. (2008) “A Canonical Model for Interactive Unawareness,” *Games and Economic Behavior* 62, 305-324.
- [11] Heifetz, A., Meier, M., Schipper, B. C. (2013) “Unawareness, Beliefs, and Speculative Trade,” *Games and Economic Behavior* 77, 100-121.
- [12] Karni, E. (2022) “A Theory-Based Decision Model,” *Journal of Economic Theory*, 201,

- [13] 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.
- [14] Karni, E., Vierø, M.-L. (2013). “Reverse Bayesianism’: A choice-based theory of growing awareness,” *American Economic Review* 103, 2790-2810.
- [15] 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.
- [16] Karni, E., Valenzuela-Stookey, Q., Vierø, M.-L. (2021). “Reverse Bayesianism’: A Generalization,” *B. E. Journal of Theoretical Economics* 21, 557-569.
- [17] Schipper, B. C. (2022) “Predicting the Unpredictable under Subjective Expected Utility” unpublished manuscript.
- [18] 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.
- [19] Vierø, M.-L. (2021) “An Intertemporal Model of Growing Awareness,” *Journal of Economic Theory* 197, 105351.
- [20] Zabell, S. L. (1992) “Predicting the Unpredictable,” *Synthese* 90, 205-232.