An intertemporal model of growing awareness

Marie-Louise Vierø*
Queen's University

February 17, 2021

Abstract

This paper presents an intertemporal model of growing awareness. It provides a framework for analyzing problems with long time horizons in the presence of growing awareness and awareness of unawareness. The framework generalizes both the standard event-tree framework and the framework from Karni and Vierø (2017) of awareness of unawareness. Axioms and a representation are provided along with a recursive formulation of intertemporal utility. This allows for tractable and consistent analysis of intertemporal problems with unawareness. To illustrate the relevance of growing awareness for dynamic decision making, the model is applied to a representative agent intertemporal asset pricing problem.

Keywords: Awareness, Unawareness, Reverse Bayesianism, Intertemporal Utility, Recursive Utility, Asset Pricing
JEL classification: D8, D81, D83

*Queen's University, Department of Economics, Dunning Hall Room 231, 94 University Avenue, Kingston, Ontario K7L 3N6, Canada. Email: viero@econ.queensu.ca. Financial support from SSHRC grant number 435-2014-0642 is gratefully acknowledged. I thank Scott Condie, Adam Dominiak, Simon Grant, Edi Karni, Yiannis Kipouros, Morten Nielsen, John Quiggin, Anna Rubinchik, Gerelt Tserenjigmid and seminar participants at Brigham Young University, Copenhagen Business School, Johns Hopkins University, Sacred Heart University, University of Maryland, University of Melbourne, University of Technology Sydney, Econometric Society World Congress 2020, RUD conference 2018, SAET conference 2018, Australian National University Economics and Philosophy Conference, the Conference on Unawareness and Unanticipated Consequences, the Decisions, Markets, and Networks conference in honour of Larry Blume and David Easley, and the Severe Uncertainty and Climate Policy Conference at the Institute for Future Studies for comments and suggestions.

1 Introduction

Under the Bayesian paradigm, the state space is fixed. As new information becomes available, the universe shrinks as some states become null. However, there are many situations in which our universe in fact expands as we become aware of new opportunities. That is, there are, quoting United States former Secretary of Defense Donald Rumsfeld, "unknown unknowns" that we may learn about.[1] In other words, a decision maker's awareness may grow over time, and the decision maker may be aware of this possibility.

Examples of new discoveries include the discovery of new diseases such as COVID-19, HIV, or mad cow disease. These are examples where familiar actions lead to discoveries of new consequences and, for many, the realization that familiar actions may continue to give rise to these, or other, new consequences. The COVID-19 pandemic has taught people to "expect the unexpected". The responses of decision makers' expectations to new and unfamiliar discoveries during the pandemic and, for example, the resulting effect on the evolution of stock prices, highlight the importance of incorporating unawareness and cognizance thereof in a dynamic context. Growing awareness has implications that reach beyond the immediate future.

This paper provides a tractable and axiomatically founded intertemporal model that incorporates growing awareness and awareness of unawareness into problems with long time horizons. The framework operationalizes the analysis of, for example, many macro and finance problems where unforeseen surprises are likely to be of importance. New issues arise in a dynamic context, because new discoveries change the outlook to the future. As an illustration, the model is applied to an intertemporal asset pricing problem, showing that unforeseen surprises and awareness of these are indeed of importance, since growing awareness in a long-horizon model affects the formation and evolution of expectations.

The analysis builds on the reverse Bayesianism framework of Karni and Vierø (2013, 2015, 2017). However, these papers considered a one-shot increase in a decision maker's awareness. They provided a framework for analyzing such an increase and axiomatized the decision maker's choice behavior in response to the increased awareness. In Karni and Vierø (2013, 2015) the decision maker is myopic with respect to her own unawareness and never anticipates making future discoveries. In Karni and Vierø (2017), the decision maker is aware of her unawareness, so although she cannot know exactly what she is unaware of, she is aware that there may be aspects of the universe that she cannot describe with her current language.

When an agent looks forward over many future periods, she can envision a plethora of ways that her awareness may grow over time. At each point in time, there are not only the

[1] U.S. Department of Defense news briefing February 12, 2002.

possibilities of making a new discovery or not, but also the possibility of making multiple new discoveries at the same time and different numbers of possible simultaneous discoveries. Thus, the possible paths of resolutions of uncertainty are much more complicated than in a standard event tree. To stay with the tree analogy, under growing awareness branches can sprout in many places in the event tree, and there will be different sprouts, and a different number of sprouts, on different branches. One challenge is to provide an analytical framework that captures all the aspects of the problem described above while at the same time keeping the problem tractable. This is achieved with a generalized event-tree in which histories reflect the evolution of both awareness and known uncertainties. The generalized event tree has the standard event tree and the state spaces in Karni and Vierø (2017) as special cases.

At a conceptual level, the novelty is to be found in the interaction of intertemporal decision-making and awareness of unawareness. The main problem that unawareness and the decision-maker's cognizance thereof poses in this setting compared to the traditional one is the impossibility of designing an optimal contingent plan for how to proceed from every decision node down the tree that remains optimal once the node is reached. When there is awareness of unawareness, the problem requires replacement of optimization over a set of well-defined contingent plans by some alternative. Such alternative is provided in the present paper.

Elaborating on the above, in the intertemporal context it is not possible to commit to a future course of action. An important question to answer, therefore, is whether unawareness and the decision-maker's cognizance thereof affect if and how the decision-maker will indeed change her course of action. Of at least as great importance is whether awareness of unawareness changes her *expectations* about her future course of action in response to awareness growth, and how these expectations will evolve in response to partial resolutions of uncertainty, i.e. to the decision-maker gradually becoming more aware. Given the important role of expectations in determining economic outcomes, understanding how growing awareness impacts the formation and evolution of expectations is important.

To give a preview of the answers to these questions, preferences depend on what the decision-maker is aware of. Therefore, the utility that a decision-maker derives from a particular choice depends not only on what she chooses, but also on what she is aware of. With a long time-horizon, the latter dependence is of particular importance, because new discoveries change the outlook to the future. Awareness growth therefore affects both the decision-maker's current choice behavior and her expectations about the future. The exact nature of this dependence, the evolution thereof, as well as its effects on expectations formations are all conceptually new insights arising from moving to the long time-horizon.

In general, future acts are not fully describable with respect to current awareness. If awareness grows in the future, the decision maker will then know a larger set of consequences than she can currently describe. She does not yet know the utility she will derive from these currently unknown consequences, and different such consequences are indistinguishable given the current level of awareness. In order to formulate preferences from the decision maker's current point of view, these issues must be dealt with by the axiomatic structure.

The axioms belong to two groups. One group consists of standard axioms. These are roughly the subjective expected utility axioms and the common intertemporal axioms of consequentialism and dynamic consistency. However, due to growing awareness, it is only reasonable to impose dynamic consistency looking forward.

The second group consists of axioms that are conceptually new and relate specifically to unawareness in an intertemporal context. In the long-horizon framework, the decision-maker anticipates that awareness may gradually grow over time. This raises the need for axioms that address how her expectations about her unawareness evolve over time. One of the new axioms serves the purpose of "preventing the agent's head from exploding". In somewhat more scientific language, the axiom assumes that the decision maker acts as if she simplifies the universe by "collapsing" unknown consequences in a particular way. Another new axiom concerns the decision maker's attitude towards yet undiscovered consequences and the evolution of this attitude as awareness grows.

The main result is an intertemporal representation of preferences. At any point in time, the agent can make contingent plans, also for events that involve new discoveries, to the extent that she can describe these plans. The axiomatic structure ensures dynamic consistency looking forward, but not necessarily looking backwards. When awareness grows, the agent may wish to change her course of action in response to her new awareness. She will, however, still maintain that her original plan was the right one given the awareness she had at the time it was made. Thus, the agent is rational to the extent possible given her limited awareness.

A recursive formulation of the decision-maker's utility is also obtained. However, the decision-maker is aware that her utility function may change in the future in response to increased awareness. Consequently, she uses an estimate, or forecast, of her future utility function, based on her current awareness, in the recursive formulation. This recursive formulation makes possible convenient analysis of, and accommodation of awareness and growing awareness in, a large class of dynamic problems.

To illustrate the relevance of growing awareness for dynamic decision making, the model is applied to a representative agent intertemporal asset pricing problem. There is one productive asset that yields a dividend each period. In the present context, the dividend should be understood abstractly as the outcome of production in the surrounding business environment, not literally as a numerical value. At each point in time, the agent entertains the possibility that she may encounter expansions of her awareness in the form of dividend realizations that are not currently known, and such expansions may indeed occur. Each period, the consumer must decide how to allocate her wealth, deriving from her asset holdings, between consumption and the asset. Let $\mathbf{c}^{est(\omega_t)} \equiv \{c_{\omega_\tau}^{est(\omega_t)}\}$ denote the consumer's estimate, when standing at history ω_t, of the process of her future consumptions. The axioms and representation imply that the consumer maximizes utility of the following form:

$$V_{\omega_t}(\mathbf{c}^{est(\omega_t)}) = u(c_{\omega_t}) + E_{\omega_t}\left[\sum_{\tau=t+1}^{T} \beta^{\tau-t} u(c_t^{est(\omega_t)})\right], \tag{1}$$

subject to anticipated budget constraints for each future history that also use estimated values of future consumptions as well as of asset holdings. In (1), E_{ω_t} denotes the agent's expectation at history ω_t, given the information and awareness she has at that history. As usual, β is the discount factor and $u(\cdot)$ is a utility function. The problem describes a consumer who is acting optimally, and as consistently as possible, given her limited awareness. She has to rely on estimates, which are her best guesses given her current awareness, when making plans for her current and future allocations of her budget. If new dividends are discovered, her actual future choices will differ from planned, because her preferences change in response to changes in her awareness. In particular, the agent's anticipation of future marginal trade-offs change, which in turn changes the trade-off between now and the future.

The structure of the utility function in (1), as well as rules for how estimates and expectations are formed and evolve as uncertainty resolves and awareness expands, are derived from the main theoretical results in this paper. Closing the model with equilibrium conditions and solving it results in processes of equilibrium prices and equilibrium price expectations. Due to growing awareness, expectations may turn out to be wrong, and both prices and price expectations fluctuate whenever there are changes in awareness. The result is an equilibrium effect of growing awareness: Asset prices as well as expectations thereof are more volatile than without growing awareness. The application illustrates one of the conceptual benefits of moving to a long time-horizon, namely that it allows for the analysis of the impact of growing awareness on the price expectation process, in particular how and when fluctuations in expectations occur in relation to when awareness grows. To address how expectations about future prices evolve in response to resolution of uncertainty, at least three periods are needed.

The paper is organized as follows: Section 2 discusses related literature. Section 3 presents the framework for modeling long time horizon problems with awareness of unawareness. Section 4 presents and discusses the axioms, while Section 5 contains the representation results. Section 6 contains the application of the model to the intertemporal asset pricing problem, and Section 7 concludes. The proofs of the main results are in the appendix.

2 Related literature

Koopmans (1972) provides an axiomatization of intertemporal utility in a deterministic framework. Epstein and Schneider (2003) axiomatize an intertemporal version of multiple-priors utility. As is the case in the present paper, they impose axioms on the entire preference process, i.e. on conditional preferences at each time-event pair. They also connect preferences conditional on different time-event pairs, rather than simply applying their axioms to conditional preferences at each time-event pair separately. The approach taken in the present paper of specifying acts from the start to the end of the event tree is inspired by Epstein and Schneider's model.

In the statistical literature, Walley (1996) and Zabell (1992) have considered related problems. Neither approach is choice theoretic. Halpern, Rong, and Saxena (2010) consider Markov decision problems with unawareness. Their decision maker is initially aware of only a subset of states and actions, and their model provides a special explore action by playing which the decision maker may become aware of actions he was previously unaware of. Halpern et al. provide conditions under which the decision maker can learn to play near-optimally in polynomial time.

Easley and Rustichini (1999) consider a decision maker who must repeatedly choose an action from a finite set. The decision maker knows the set of available actions and that a payoff will occur to each action in each period, but no further structure. The decision maker prefers more payoff to less. He begins with an arbitrary ordering over acts and selects the action with the highest rank. Upon resolution of the period's uncertainty, he observes the payoff to each action and updates his ordering. Easley and Rustichini provide axioms that lead to actions eventually being chosen optimally according to expected utility.

Grant, Meneghel, and Tourky (2017) present a model of learning after an expansion in awareness. When awareness grows, the state space expands. Following the expansion, beliefs are initially imprecise, but over time they converge to a precise probability distribution.

Grant and Quiggin (2013a, 2013b) consider dynamic games with differential awareness, where players may be unaware of some histories of the game. Unawareness thus materializes as players considering only a restricted version of the game. For such games, Grant and Quiggin provide logical foundations for players using inductive reasoning to conclude that there may be propositions, and hence parts of the game tree, of which they are unaware. Players may also gain inductive support for particular actions leading to unforeseen contingencies. Ozbay (2007) considers games where a player's awareness may increase due to strategic announcements by his opponent.

There are a number of papers taking a choice theoretic approach to unawareness or related issues. These include Li (2008), Ahn and Ergin (2010), Schipper (2013), Lehrer and Teper (2014), Kochov (2018), Walker and Dietz (2011), Alon (2015), Grant and Quiggin (2015), Dietrich (2018), Piermont (2017), and Dominiak and Tserenjigmid (2017). Kochov (2018) uses a three-period model to distinguish between unforeseen and ambiguous events. The other papers are either static in nature or consider one-shot increases in awareness.

Since the present paper builds on Karni and Vierø (2013, 2015, 2017), it is useful to describe these works in somewhat more detail. Karni and Vierø (2013) considers a one-shot increase in a decision maker's awareness. There are two main contributions. The first is to provide a framework of an expanding universe. What Karni and Vierø call the conceivable state space expands as new acts, consequences, or links between them are discovered, that is, when awareness grows. The second contribution is to invoke the revealed preference methodology and axiomatize the decision maker's choice behaviour in the expanding universe. The challenge is that preferences under different levels of awareness are defined over different domains, so they need to be linked. The axioms imply that for a given level of awareness, the decision maker is an expected utility maximizer. The axioms that link behaviour across different state spaces imply that the utility of known risks is invariant to expansions of awareness and also imply reverse Bayesian updating of beliefs: when new discoveries are made, probability mass is shifted proportionally away from events in the prior state space to events created as a result of the expansion of the state space.

Karni and Vierø (2015) has a more general preference structure within the same framework. In both Karni and Vierø (2013) and Karni and Vierø (2015) the decision maker is myopic with respect to her unawareness. Hence, she never anticipates making future discoveries and always acts as if she is fully aware.[2]

The premise of Karni and Vierø (2017) is that if you have become aware of new things in the past, you may anticipate that this can also happen in the future. The paper also considers only a one-shot increase in the decision maker's awareness and extends the framework from Karni and Vierø (2013) to allow for decision makers being aware of their unawareness. So, although decision makers cannot know exactly what they are unaware of, they are aware that there may be aspects of the universe that they cannot describe with their current language. The paper limits attention to growing awareness of consequences. The framework has an augmented conceivable state space which is partitioned into fully describable and imperfectly describable states, in the latter of which awareness expands. The axiomatic structure implies that for a given level of awareness, the decision maker is a generalized Expected Utility maximizer: the utility representation consists of a Bernoulli utility function over known consequences, beliefs over the augmented conceivable state space that assign

[2]Karni, Valenzuela-Stookey and Vierø (forthcoming) generalize the result in Karni and Vierø (2013).

beliefs to expansions in the decision maker's awareness, and an extra parameter that reflects the decision maker's attitude towards the unknown. As in Karni and Vierø (2013), there is reverse Bayesian updating of beliefs and the utility of known risks is invariant to expansions in awareness. However it is now also possible to characterize the decision maker's sense of ignorance, which is captured by the probability assigned to expansions in her awareness, and the evolution thereof.

3 Analytical Framework

Time is discrete, indexed by $t \in \{0, 1, \ldots, T\}$, where T is finite. The decision maker is aware of this. Let the initial state of the world, which is known by the decision maker, be denoted by s_0. Let A be a finite, nonempty, set of basic actions with generic element a. The set of basic actions is available in each period, known by the decision maker, and remains fixed throughout. In contrast, the set of known feasible consequences evolves over time as the decision maker's awareness grows.[3] Let $C(s_0)$ be the initial set of known feasible consequences, which is finite and nonempty. The elements of $C(s_0)$ are the consequences the decision maker is initially aware of.

For any set of consequences C, let c denote a generic element and define $x(C) = \neg C$ to be the abstract "consequence" that has the interpretation "none of the above" and captures consequences of which the decision maker is currently unaware. There may, in fact, be any number of unknown consequences or no such consequence at all. The consequence $x(C)$ includes all of these possibilities. From an ex ante perspective, the decision maker cannot know whether $x(C)$ is a singleton, a non-degenerate set, or an empty set. Define $\widehat{C} = C \cup \{x(C)\}$, referred to as the set of extended consequences, with generic element $\hat{c}$. Label by $\hat{c}^1, \hat{c}^2, \ldots$ the currently unknown consequences in order of discovery.

From a time-0 perspective, the only well-defined consequences are those in $C(s_0)$ and $x(C(s_0))$. From a time-0 perspective any yet undiscovered consequences $\hat{c}^1, \hat{c}^2, \ldots$ are all "none-of-the-above" and thus part of, or indistinguishable from, $x(C(s_0))$ and also indistinguishable from each other. However, the decision maker does know that when she has to make future choices, she may have discovered additional consequences. If her awareness grows, she will be able to distinguish consequences that are currently indistinguishable.

[3]The present paper thus follows Karni and Vierø (2017) in considering growing awareness of consequences while keeping awareness of basic actions constant. Karni and Vierø (2013) shows that in a one-shot framework without awareness of unawareness, discovery of new actions materializes differently in the evolution of the state space than the discovery of new consequences. The question of awareness of unawareness of actions has not yet been addressed in a one-shot framework. Given this, it is beyond the scope of the present paper.

Therefore, her ex-ante and ex-post views of a problem are different under growing awareness.

Below, state spaces, which depict possible one-step-ahead resolutions of uncertainty and awareness, as well as the space of possible histories are defined. For notational convenience, let $x_s \equiv x(C(s))$ for any state s.

3.1 State spaces

A state space depicts the possible one-step-ahead resolutions of uncertainty. Define the time-1 state space by

$$S_1(s_0) \equiv (\widehat{C}(s_0))^A = \{s : A \to \widehat{C}(s_0)\},$$

which is the set of all functions from the set of basic actions to the initial set of extended consequences. Hence, a state specifies the unique extended consequence that is associated with each of the basic actions. The time-1 state space thus depicts the possible resolutions of uncertainty at $t = 1$. This object was referred to as the augmented conceivable state space in Karni and Vierø (2017). It exhausts all the possible ways one can assign extended consequences to the basic actions. Define also the set

$$\tilde{S}_1(s_0) \equiv (C(s_0))^A = \{s : A \to C(s_0)\},$$

i.e. the set of functions from basic actions to the initial set of known consequences. This is the subset of $S_1(s_0)$ containing fully describable states whose description only involves known consequences. The complement $S_1(s_0) \setminus \tilde{S}_1(s_0)$ is referred to as the set of imperfectly describable states, since their descriptions include the unknown consequence x_{s_0}, which only has an abstract meaning. A generic time-1 state is denoted by s_1. Example 1 provides an illustration.

Example 1 Consider the situation in which there are two basic actions, $A = \{a_1, a_2\}$, and two initially known feasible consequences, $C(s_0) = \{c_1, c_2\}$. The none-of-the-above consequence x_{s_0} is any consequence different from c_1 and c_2 that potentially could be discovered. The time-1 state space $S_1(s_0)$ consists of the nine states depicted in the following matrix:

$$\begin{array}{cccccccccc} & s_1^1 & s_1^2 & s_1^3 & s_1^4 & s_1^5 & s_1^6 & s_1^7 & s_1^8 & s_1^9 \\ a_1 & c_1 & c_2 & c_1 & c_2 & x_{s_0} & x_{s_0} & c_1 & c_2 & x_{s_0} \\ a_2 & c_1 & c_1 & c_2 & c_2 & c_1 & c_2 & x_{s_0} & x_{s_0} & x_{s_0} \end{array} \qquad (2)$$

The subset of fully describable states is $\tilde{S}_1(s_0) = \{s_1^1, \ldots, s_1^4\}$, while $S_1(s_0) \setminus \tilde{S}_1(s_0) = \{s_1^5, \ldots, s_1^9\}$ are imperfectly describable.■

As it appears from Example 1 and matrix (2), the time-1 states in $S_1(s_0)$ differ in how many previously unknown consequences will be discovered. In each of the fully describable states $s_1^1, \ldots, s_1^4$, no new consequence is discovered. In each imperfectly describable state, new consequences are discovered. One new consequence is discovered in each of $s_1^5, \ldots, s_1^8$, and two potentially different new consequences are discovered in state s_1^9. The set of known feasible consequences that the decision maker is aware of at time 1 thus depends on what is discovered at time 1, i.e. it is a function of which state is realized in the first period.

Define $n(s_1)$ as the number of previously unknown consequences discovered in s_1. It is assumed that when a basic action reveals something previously unknown at a particular point in time, whatever it is, it is considered as one consequence. Thus, $n(s_1) \in \{0, \ldots, |A|\}$. Let $\{\hat{c}^i(s_1)\}_{i=1}^{n(s_1)}$ be the set of new consequences discovered in s_1, with $\{\hat{c}^i(s_1)\}_{i=1}^{n(s_1)} \equiv \emptyset$ if $n(s_1) = 0$. Then the set of known feasible consequences at time 1 is given by

$$C(s_1) \equiv C(s_0) \cup \{\hat{c}^i(s_1)\}_{i=1}^{n(s_1)},$$

while the abstract "none-of-the-above" consequence is $x_{s_1} = \neg C(s_1)$, and the set of extended consequences is $\widehat{C}(s_1) = C(s_1) \cup \{x_{s_1}\}$. Similar to the definition of the time-1 state space, define the time-2 state space originating at state s_1 by $S_2(s_1) \equiv (\widehat{C}(s_1))^A$. That is, $S_2(s_1)$ depicts the possible one-step-ahead resolutions of uncertainty following s_1. Also define the subset of fully describable time-2 states $\tilde{S}_2(s_1) \equiv (C(s_1))^A$ and denote a generic element by s_2.

Example 2 Consider a situation with two basic actions, $A = \{a_1, a_2\}$, and initially just one known feasible consequence, $C(s_0) = \{c_1\}$. Then the none-of-the-above consequence x_{s_0} is any consequence different from c_1. The time-1 state space thus consists of the four states in matrix (3) below:

$$\begin{array}{ccccc} & s_1^1 & s_1^2 & s_1^3 & s_1^4 \\ a_1 & c_1 & c_1 & x_{s_0} & x_{s_0} \\ a_2 & c_1 & x_{s_0} & c_1 & x_{s_0} \end{array} \tag{3}$$

In the fully describable state s_1^1, no new consequence is discovered. Hence, $C(s_1^1) = C(s_0)$. As a result, $x_{s_1^1} = x_{s_0}$, and $S_2(s_1^1) = S_1(s_0)$, i.e. as depicted in (3).

In the imperfectly describable state s_1^2, one new consequence, $\hat{c}^1(s_1^2)$, is discovered. Therefore, $C(s_1^2) = C(s_0) \cup \{\hat{c}^1(s_1^2)\}$. Then $x_{s_1^2} = \neg\{c_1, \hat{c}^1(s_1^2)\}$, the set of extended consequences is $\widehat{C}(s_1^2) = \{c_1, \hat{c}^1(s_1^2), x_{s_1^2}\}$ and the time-2 state space following state s_1^2, $S_2(s_1^2) = (\widehat{C}(s_1^2))^A$, consists of 9 states as depicted in matrix (4) below, where $\hat{c}^1 \equiv \hat{c}^1(s_1^2)$:

$$
\begin{array}{cccccccccc}
 & s_2^1 & s_2^2 & s_2^3 & s_2^4 & s_2^5 & s_2^6 & s_2^7 & s_2^8 & s_2^9 \\
a_1 & c_1 & c_1 & \hat{c}_1 & \hat{c}_1 & c_1 & \hat{c}_1 & x_{s_1^2} & x_{s_1^2} & x_{s_1^2} \\
a_2 & c_1 & \hat{c}_1 & c_1 & \hat{c}_1 & x_{s_1^2} & x_{s_1^2} & c_1 & \hat{c}_1 & x_{s_1^2}
\end{array} \tag{4}
$$

The situation if s_1^3 is realized is similar to that if s_1^2 is realized, except that the consequence $\hat{c}^1(s_1^3)$ that is discovered in s_1^3 could be different from that which would be discovered if s_1^2 were realized. Since $\hat{c}^1(s_1^3)$ is potentially different from $\hat{c}^1(s_1^2)$, the sets $C(s_1^2)$ and $C(s_1^3)$ are potentially different. Consequently, $x_{s_1^3}$ and $x_{s_1^2}$ are potentially different, as are $S_2(s_1^3)$ and $S_2(s_1^2)$. Importantly, from an ex-ante perspective, the decision maker cannot distinguish between different unknown consequences, since she is unaware of their attributes. However, she can reason, like we just did, that they can potentially be different. The decision maker can envision that the situation following s_1^3 may be different than that following s_1^2. The time-2 state space following state s_1^3, $S_2(s_1^3)$, is as depicted in (4), with $x_{s_1^2}$ replaced by $x_{s_1^3}$ and $\hat{c}^1$ appropriately redefined.

In s_1^4, two new consequences $\hat{c}^1(s_1^4)$ and $\hat{c}^2(s_1^4)$ are discovered. It could be that $\hat{c}^1(s_1^4) = \hat{c}^2(s_1^4)$, but from an ex-ante perspective using distinct $\hat{c}^1(s_1^4)$ and $\hat{c}^2(s_1^4)$ allows the decision maker to formulate the maximal increase in awareness that she can anticipate. Then $C(s_1^4) = C(s_0) \cup \{\hat{c}^1(s_1^4), \hat{c}^2(s_1^4)\}$ and $x(C(s_1^4)) = \neg\{c_1, \hat{c}^1(s_1^4), \hat{c}^2(s_1^4)\}$. The set of extended consequences is $\widehat{C}(s_1^4) = \{c_1, \hat{c}^1(s_1^4), \hat{c}^2(s_1^4), x_{s_1^4}\}$, and the time-2 state space following state s_1^4, $S_2(s_1^4) = (\widehat{C}(s_1^4))^A$, consists of 16 elements as in matrix (5), where $(\hat{c}^1, \hat{c}^2) = (\hat{c}^1(s_1^4), \hat{c}^2(s_1^4))$:

$$
\begin{array}{ccccccccccccccccc}
 & s_2^1 & s_2^2 & s_2^3 & s_2^4 & s_2^5 & s_2^6 & s_2^7 & s_2^8 & s_2^9 & s_2^{10} & s_2^{11} & s_2^{12} & s_2^{13} & s_2^{14} & s_2^{15} & s_2^{16} \\
a_1 & c_1 & c_1 & c_1 & \hat{c}_1 & \hat{c}_1 & \hat{c}_1 & \hat{c}_2 & \hat{c}_2 & \hat{c}_2 & c_1 & \hat{c}_1 & \hat{c}_2 & x_{s_1^4} & x_{s_1^4} & x_{s_1^4} & x_{s_1^4} \\
a_2 & c_1 & \hat{c}_1 & \hat{c}_2 & c_1 & \hat{c}_1 & \hat{c}_2 & c_1 & \hat{c}_1 & \hat{c}_2 & x_{s_1^4} & x_{s_1^4} & x_{s_1^4} & c_1 & \hat{c}_1 & \hat{c}_2 & x_{s_1^4}
\end{array} \tag{5}
$$

From the time-0 perspective, the total number of time-2 states is 4+9+9+16=38.■

In general, for $t > 0$, let s_t denote a generic state and define $n(s_t)$ to be the number of previously unknown consequences discovered in s_t. Let $\{\hat{c}^i(s_t)\}_{i=1}^{n(s_t)}$ be the set of new consequences discovered in s_t, with $\{\hat{c}^i(s_t)\}_{i=1}^{n(s_t)} \equiv \emptyset$ if $n(s_t) = 0$. Then the set of known feasible consequences in s_t is given by

$$C(s_t) \equiv C(s_{t-1}) \cup \{\hat{c}^i(s_t)\}_{i=1}^{n(s_t)},$$

while the abstract "none-of-the-above" consequence is $x_{s_t} = \neg C(s_t)$, and the set of extended consequences is $\widehat{C}(s_t) = C(s_t) \cup \{x_{s_t}\}$.

Define the time $t+1$ state space originating in state s_t by

$$S_{t+1}(s_t) \equiv (\widehat{C}(s_t))^A,$$

which depicts the possible one-step-ahead resolutions of uncertainty following s_t. Define also the subset of $S_{t+1}(s_t)$ containing fully describable time $t+1$ states originating in state s_t by

$$\tilde{S}_{t+1}(s_t) \equiv (C(s_t))^A.$$

3.2 The history space

The history space can be depicted by an event tree, albeit non-standard. Define, for any $t \in \{0, \dots, T\}$, the time-t *partial* history space to be the set

$$\Omega_t \equiv \{\omega_t = (s_0, s_1, \dots, s_t) : \; s_\tau \in S_\tau(s_{\tau-1}) \; \forall \, \tau = 1, \dots, t\}, \tag{6}$$

with representative element $\omega_t \equiv (s_0, s_1, \dots, s_t)$. Thus, a time-$t$ partial history is a complete description of the resolution of uncertainty and evolution of the decision maker's awareness up to and including time t. For notational convenience, let $\Omega \equiv \Omega_T$ denote the full history space, that is, the set of all possible evolutions of uncertainty and the decision maker's awareness at all times. Thus, a full history $\omega \equiv \omega_T$ is a complete description of the resolution of uncertainty and evolution of the decision maker's awareness at all times, that is, a complete path through the event tree. Histories differ from an ex-ante and an ex-post perspective. From an ex-ante perspective, potential future histories that involve increases in awareness can only be abstractly described as involving yet undiscovered consequences. From an ex-post perspective, in a particular history, the discovered consequences are now known, so the history can be concretely described. However, alternative histories that did not materialize can still only be abstractly described. Let

$$\mathbf{P}_\tau(\omega) \equiv (s_0, s_1, \dots, s_\tau) \tag{7}$$

denote the projection of history ω onto Ω_τ, that is, the first $\tau + 1$ components of history ω. Thus, $\mathbf{P}_\tau(\omega) \in \Omega_\tau$.

Define, for all $t \in \{0, \dots, T\}$,

$$\boldsymbol{\Omega}_t \equiv \bigcup_{\tau=0}^{t} \Omega_\tau. \tag{8}$$

This is the set of all partial histories that end at or before time t.[4] For ease of notation, let $\boldsymbol{\Omega} \equiv \boldsymbol{\Omega}_T$. Hence, $\boldsymbol{\Omega}$ is the set of all partial and complete paths through the event tree. In other words, it is the set of all nodes in the event tree, where branches start or end.

[4] Notice the bold font.

Example 2 (continued) The event tree for the situation with $A = \{a_1, a_2\}$, $C(s_0) = \{c_1\}$, and T=3 is depicted in Figure 1. The numbers after each time-2 partial history indicate the number of branches originating at that partial history, and thus give the number of possible time-3 states originating at that time-2 partial history. The number of states and histories, and hence the possible evolutions of awareness, quickly becomes very large. There are 4 time-1 partial histories, 38 time-2 partial histories, and 618 time-3 partial histories. In a standard model with 4 time-1 histories, there would be 16 time-2 histories and 64 time-3 histories.■

Note that while n, C, S, etcetera are defined recursively one step ahead as functions of s_t, they can also be described as functions of the partial history ω_t: $n(\omega_t)$, $C(\omega_t)$, x_{ω_t}, $S_{t+1}(\omega_t)$, and $\tilde{S}_{t+1}(\omega_t)$. Define, for each $\omega_t \in \mathbf{\Omega}$ and $\tau \in \{t, \dots, T\}$,

$$\Omega_\tau(\omega_t) \equiv \{\omega_\tau = (\omega_t, s_{t+1}, \dots, s_\tau) : s_{t+1} \in S_{t+1}(\omega_t) \text{ and } s_{\hat{t}} \in S_{\hat{t}}(s_{\hat{t}-1}) \; \forall \; \hat{t} = t+2, \dots, \tau\}. \quad (9)$$

The set in (9) is the set of time-τ partial histories that can be reached from history ω_t, or, in other words, the set of possible continuation paths through time τ, starting from partial history ω_t. Then the set of full histories that can be reached from ω_t is $\Omega(\omega_t) \equiv \Omega_T(\omega_t)$.

For each $t \in \{0, \dots, T\}$, let $\mathbf{I}_t \equiv \{\Omega(\omega_t) : \omega_t \in \Omega_t\}$, and note that $\mathbf{I}_t$ forms a finite partition of Ω. Let $\mathscr{I}_t$ be the σ-algebra generated by the partition $\mathbf{I}_t$. Then $\mathscr{I} = \{\mathscr{I}_t\}_{t=0}^T$ is an increasing sequence of σ-algebras, i.e. a filtration. The filtration $\mathscr{I}$ represents the information structure, including the decision maker's awareness, with the caveat that consequences to be discovered in the future are only abstractly described. In plain English, the decision maker knows the structure of the event tree, where in the event tree she is, and the concrete nature of consequences that have been discovered at the partial history at which she finds herself.

Finally, parallel to the definition in (8), define

$$\mathbf{\Omega}_\tau(\omega_t) \equiv \bigcup_{i=t}^{\tau} \Omega_i(\omega_t), , \quad (10)$$

which is the set of all continuation paths or continuation histories reachable from ω_t ending at or before time τ, including ω_t itself. Let $\mathbf{\Omega}(\omega_t) \equiv \mathbf{\Omega}_T(\omega_t)$.

Many of these possible future partial and full histories, as well as the consequences the decision maker can obtain in them, are indescribable beyond "there may be a number of currently unknown consequences that I could potentially have discovered by then" from her current point of view. Her current point of view is reflected by ω_t. In other words, the future is not fully describable with respect to the decision maker's current awareness.

The framework introduced above captures the important aspects of the problem of awareness of unawareness with long time horizons, namely that there is a plethora of ways that

Figure 1: History space for Example 2. The numbers after each time-2 partial history indicate the number of branches originating at that partial history.

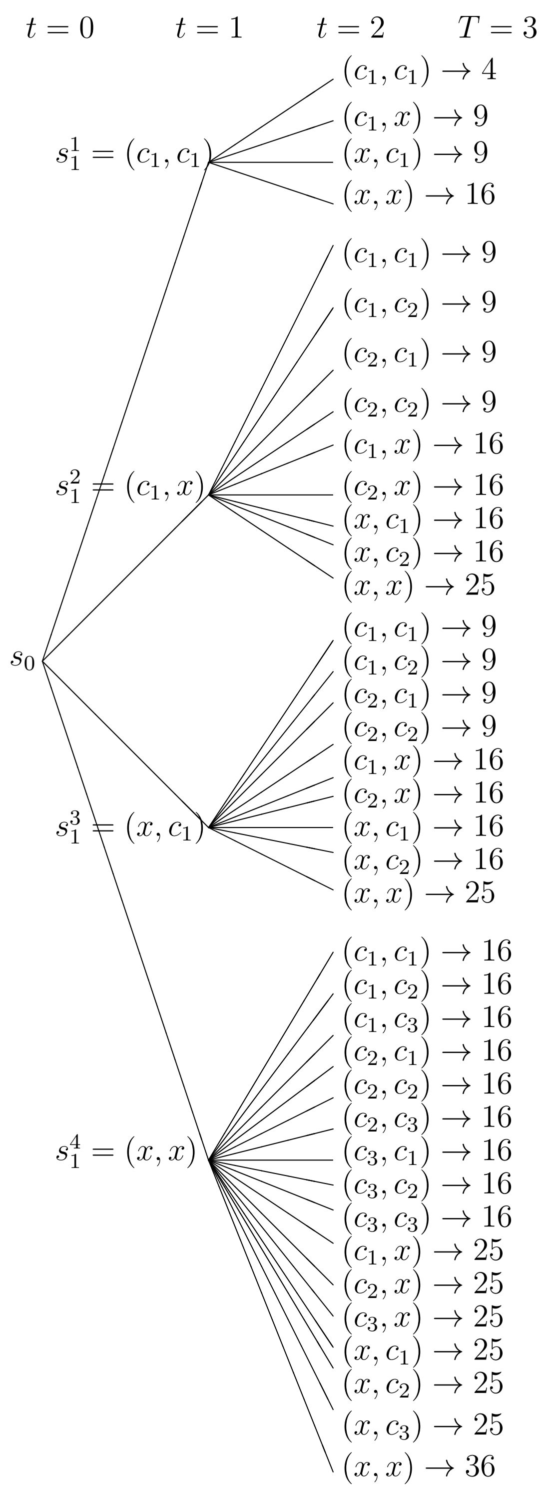

awareness can evolve both in terms of when, what, how much, and in which order discoveries are made. The framework does so in a systematic way that nests the standard approach of using event trees. It also nests the state spaces in Karni and Vierø (2017). It is the simplest generalization that nests both models.

3.3 Conceivable intertemporal acts

Since this paper uses the revealed preference methodology, it is a requirement that, for a given level of awareness, bets can be both meaningfully described using current language and settled once uncertainty has been resolved. For each non-ultimate partial history $\omega_t \in \mathbf{\Omega}_{T-1}$, define

$$f(\omega_t) : S_{t+1}(\omega_t) \to \Delta(\widehat{C}(\omega_t)) \text{ such that } f(\omega_t)(s) \in \Delta(C(\omega_t)) \text{ for all } s \in \tilde{S}_{t+1}(\omega_t), \quad (11)$$

where $\Delta(\cdot)$ denotes the probability simplex.[5] I.e. $f(\omega_t)$ is a function from $S_{t+1}(\omega_t)$ into the set of lotteries over the set of extended consequences $\widehat{C}(\omega_t)$ that the decision maker is aware of at partial history ω_t. The supports of the lotteries are restricted to the set of known consequences $C(\omega_t)$ in the ω_t-fully describable states in $S_{t+1}(\omega_t)$. See Figure 2 for an illustration. The act defined in (11) is referred to as a restricted Anscombe-Aumann act originating at partial history ω_t. It is a one-step-ahead act in the sense that the uncertainty regarding it will be resolved at the beginning of the next period. In (11), the notation ω_t is used to denote the originating partial history and s is used to denote the next period states in which the payoff of the one-step-ahead restricted Anscombe-Aumann act materializes.

The reason for the range being restricted in the fully describable states is the requirement that bets should be possible to settle once uncertainty resolves, and that decision makers cannot meaningfully form preferences over acts that assign indescribable consequences to fully describable states. In fully describable states, the consequence x remains abstract, and one cannot deliver a consequence that has not yet been discovered. However, there is no problem with promising to deliver a consequence, which is none of the prior consequences, if such a consequence is discovered. Therefore, the acts can assign, to imperfectly describable states only, consequences that will be discovered if these states obtain.[6] As a result, the support of the lotteries in the restricted Anscombe-Aumann acts is L-shaped across states, rather than rectangular like the standard Anscombe and Aumann (1963) acts, as Figure 2 shows.

[5]The usual abuse of notation is adopted, where c is also used to denote the lottery that returns consequence c with probability 1.

[6]For further discussion of this issue, see Karni and Vierø (2017).

Figure 2: Illustration of the support of the lotteries in the restricted Anscombe-Aumann acts

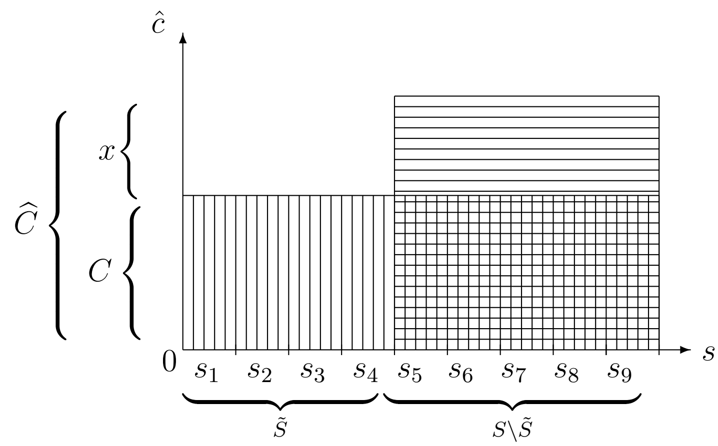

Define the set of all restricted Anscombe-Aumann acts originating at partial history ω_t:

$$F(\omega_t) \equiv \{f(\omega_t)\}.$$

This is the set of all functions described in (11). At each point in time (and in each partial history), two things happen: the uncertainty regarding the previous period's one-step-ahead act $f(\omega_{t-1})$ resolves and a new, current, one-step-ahead act $f(\omega_t)$ may be chosen. The last period differs, since no new one-step-ahead act is chosen.

Define

$$f \equiv (f(\omega_t))_{\omega_t \in \mathbf{\Omega}_{T-1}}. \tag{12}$$

The act defined in (12) is an intertemporal act, consisting of a one-step-ahead restricted Anscombe-Aumann act as defined in (11) for each partial history, that is, for each branching point in the event tree.[7] The intertemporal acts reflect that although from a time-0 perspective the only well-defined consequences are those in $C(s_0)$ and x_{s_0}, the decision maker knows that when she has to make future choices, she may have discovered additional consequences.

The set of all intertemporal acts can now be defined:

$$F \equiv \{f = (f(\omega_t))_{\omega_t \in \mathbf{\Omega}_{T-1}}\}. \tag{13}$$

This set of intertemporal acts, defined in (13), is the domain of the decision maker's preferences. It is the set of all complete contingent plans the decision maker can describe given her awareness. Thus, the decision maker has preferences over complete contingent plans, and it

[7]Here and henceforth, the terms "each state" and "each partial history" are used to refer to all states and partial histories but the ultimate-period ones. In order to keep the exposition clean, the distinction of ultimate states and histories will not be mentioned, except when it is directly relevant.

is assumed that she can change her contingent plan at each partial history, should she wish to do so.

Let $h_{\omega_t} f$ be the intertemporal act obtained from f by replacing the restricted Anscombe-Aumann act originating at ω_t by $h \in F(\omega_t)$. Also, for $E \subseteq S_{t+1}(\omega_t)$, let $\hat{p}_E f$ be the intertemporal act that returns the lottery $\hat{p}$ in all states in the event E and agrees with f elsewhere. The act $\hat{p}_E f$ is thus a special case of $h_{\omega_t} f$ for which h agrees with $f(\omega_t)$ for $s \in S_{t+1}(\omega_t) \setminus E$ and is constant at $\hat{p}$ for $s \in E$. For all $f \in F$, define

$$H_{\omega_t}(f) \equiv \{h_{\omega_t} f | h \in F(\omega_t)\},$$

which is the set of all intertemporal acts that agree with f with the exception of the restricted Anscombe-Aumann act originating at ω_t.

4 Preferences

The decision maker has a preference ordering on the set of interetemporal acts F (defined in (13)) at any partial history $\omega_t \in \mathbf{\Omega}_{T-1}$. It is denoted by $\succsim_{\omega_t}$ and expresses the ordering conditional on the awareness level prevailing given the cumulative discoveries made in partial history ω_t. Strict preference $\succ_{\omega_t}$ and indifference $\sim_{\omega_t}$ are defined as usual.[8] Axioms will be imposed on the collection of preference orderings $\{\succsim_{\omega_t} : \omega_t \in \mathbf{\Omega}_{T-1}\}$. It is henceforth assumed that $C(\omega_0)$ contains at least two elements. A history $(\omega_{T-1}, s) \in \Omega$ is said to be $\succ_{\omega_t}$-null if $\hat{p}_s f \sim_{\omega_t} \hat{q}_s f$ for all $\hat{p}, \hat{q} \in \Delta(\widehat{C}(\omega_{T-1}))$ for all $f \in F$.[9] A history is said to be $\succ_{\omega_t}$-nonnull if it is not $\succ_{\omega_t}$-null. If history $(\omega_{T-1}, s) \in \Omega$ is $\succ_{\omega_t}$-null (respectively $\succ_{\omega_t}$-nonnull), then state s is also said to be $\succ_{\omega_t}$-null (respectively $\succ_{\omega_t}$-nonnull). For any $\tau \in \{t, \ldots, T\}$, partial history ω_τ is said to be $\succ_{\omega_t}$-null if (ω_{T-1}, s) is $\succ_{\omega_t}$-null for all (ω_{T-1}, s) for which $\mathbf{P}_\tau(\omega_{T-1}, s) = \omega_\tau$. Otherwise, ω_τ is said to be $\succ_{\omega_t}$-nonnull.

4.1 Axioms

A total of nine axioms are imposed. Axioms 1 through 4 are roughly standard, albeit the statement of some of them is more tedious in the present context. Axiom 5 would be implied by the independence axiom in a standard framework. Axioms 7 and 9 impose dynamic consistency when looking forward. They extend the corresponding axioms in Karni

[8]One could argue that in order to indeed be true to the revealed preference approach, the decision maker should only be required to be able to express preferences over acts defined on the set of possible partial and full continuation histories. Under Axiom 1 below, the present definition of intertemporal acts gives an equivalent result, while staying methodologically closer to the existing literature.

[9]With the restriction on $\hat{p}, \hat{q}$ to $\Delta(C(\omega_{T-1}))$ if s is fully describable.

and Vierø (2017) to the intertemporal framework. Axioms 6 and 8 are new and, to my knowledge, cannot be found in the existing literature. They specifically relate to awareness of unawareness with a long time-horizon and the new issues that arise as a result of that.

Axioms 1 through 6 are imposed on preferences at any partial history.

Axiom 1 (Consequentialism). *For all* $\omega_t \in \mathbf{\Omega}_{T-1}$*, for all* $f, f' \in F$*, if* $f(\omega_\tau) = f'(\omega_\tau)$ *for all* $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$*, then* $f \sim_{\omega_t} f'$*.*

Axiom 1 postulates that only continuations of acts matter for preferences. Thus, at any history, the decision maker does not care about parts of the event tree that cannot be reached from her current position.

Axioms 2 through 5 resemble the axioms in Karni and Vierø (2017) that result in their generalized expected utility representation, although the present domains are different than in Karni and Vierø (2017). In particular, Axiom 2 contains the standard expected utility axioms.

Axiom 2 (Expected Utility). *For all* $\omega_t \in \mathbf{\Omega}_{T-1}$*,*

(i) (Preorder) the relation $\succ_{\omega_t}$ *is asymmetric and negatively transitive on* F*.*

(ii) (Archimedean) for all $h, h', h'' \in F$*, if* $h \succ_{\omega_t} h'$ *and* $h' \succ_{\omega_t} h''$*, then there exist* $\alpha, \beta \in (0,1)$ *such that* $\alpha h + (1-\alpha)h'' \succ_{\omega_t} h'$ *and* $h' \succ_{\omega_t} \beta h + (1-\beta)h''$*.*

(iii) (Independence) for all $h, h', h'' \in F$ *and for all* $\alpha \in (0,1]$*,* $h \succ_{\omega_t} h'$ *if and only if* $\alpha h + (1-\alpha)h'' \succ_{\omega_t} \alpha h' + (1-\alpha)h''$*.*

Axiom 3 (Monotonicity). *For all* $\omega_t \in \mathbf{\Omega}_{T-1}$*,*

(i) for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ *and* $\succ_{\omega_t}$*-nonnull* $s \in S_{\tau+1}(\omega_\tau)$*, for all* $p, q \in \Delta(C(\omega_\tau))$*, and for all* $f \in F$ *it holds that* $p_s f \succ_{\omega_t} q_s f$ *if and only if* $p_{\omega_\tau} f \succ_{\omega_t} q_{\omega_\tau} f$*.*

(ii) for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ *and* $\succ_{\omega_t}$*-nonnull* $s \in S_{\tau+1}(\omega_\tau) \backslash \tilde{S}_{\tau+1}(\omega_\tau)$*, for all* $\hat{p}, \hat{q} \in \Delta(\widehat{C}(\omega_\tau))$*, and for all* $f \in F$ *it holds that* $\hat{p}_s f \succ_{\omega_t} \hat{q}_s f$ *if and only if* $\hat{p}_{S_{\tau+1}(\omega_\tau)\backslash\tilde{S}_{\tau+1}(\omega_\tau)} f \succ_{\omega_t} \hat{q}_{S_{\tau+1}(\omega_\tau)\backslash\tilde{S}_{\tau+1}(\omega_\tau)} f$*.*

(iii) for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$*, for all* $p, q \in \Delta(C(\omega_t))$*, and for all* $f \in F$ *it holds that* $p_{\omega_\tau} f \succ_{\omega_t} q_{\omega_\tau} f$ *if and only if* $p_{\mathbf{\Omega}(\omega_t)} f \succ_{\omega_t} q_{\mathbf{\Omega}(\omega_t)} f$*.*

In Axiom 3, the content of parts (i) and (ii) are similar to the standard content of monotonicity, but the statement differs.[10] The difference in statement is necessary because the support of the lotteries in fully describable states is restricted to the set of consequences that are known in the partial history at which the restricted Anscombe-Aumann act originates, while in the imperfectly describable states, the lotteries can involve the unknown consequence that will be discovered. That is, across states the support of the lotteries is L-shaped rather than rectangular, as Figure 2 illustrates, which necessitates the statement of monotonicity as in Axiom 3. Part (iii) extends monotonicity to also hold for lotteries that occur at different points in time.

Axiom 4 (Nontriviality)**.** *For all $f \in F$, and for all $\omega_t \in \mathbf{\Omega}_{T-1}$, the strict preference relation $\succ_{\omega_t}$ is non-empty on $H_{\omega_\tau}(f)$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$.*

Axiom 4 requires non-triviality of each preference relation $\succ_{\omega_t}$ on sets of acts that only differ in the restricted Anscombe-Aumann act originating at one partial history. It implies that no partial history in the continuation path is $\succ_{\omega_t}$-null.

Axiom 5 (Separability)**.** *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $f, g \in F$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, and for all $\hat{p}, \hat{q} \in \Delta(\widehat{C}(\omega_\tau))$, it holds that $\hat{p}_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} f \succ_{\omega_t} \hat{q}_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} f$ if and only if $\hat{p}_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} g \succ_{\omega_t} \hat{q}_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} g$.*

Axiom 5 regards intertemporal acts that only differ in the restricted Anscombe-Aumann act originating in a particular future (or in the current) partial history. Furthermore, those restricted Anscombe-Aumann acts only differ on the imperfectly describable states that follow and are constant on that set of states. The Axiom requires that the ranking of such intertemporal acts is independent of the aspects on which the acts agree. This separability is not implied by the independence axiom as it would be in a standard framework, since the payoff $x(C(\omega_\tau))$ is not defined on $\tilde{S}_{\tau+1}(\omega_\tau)$.

Axiom 6 (Unknowns are Unknowns)**.** *For all $f \in F$, for all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, and for all $\hat{c} \in \widehat{C}(\omega_\tau) \setminus C(\omega_t)$, $(x_{\omega_t})_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} f \sim_{\omega_t} \hat{c}_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)} f$.*

As the name suggests, Axiom 6 requires that the decision maker treats all unknowns as such. She does not a-priori distinguish between, for example, unknowns to be discovered at different times or in different partial histories. Anything that cannot be described or imagined at her current level of awareness is treated the same way by the decision maker. This does not preclude that she will have a preference for when to make such discoveries.

[10] In parts (i) and (iii) of Axiom 3, the notation p in $p_{\omega_\tau} f$ is abused to denote the restricted Anscombe-Aumann act for which $f(\omega_\tau)(s) = p$ for all $s \in S_{\tau+1}(\omega_\tau)$. The act $p_{\mathbf{\Omega}(\omega_t)} f$ returns p everywhere in the continuation path.

When setting $\hat{c} = x_{\omega_\tau}$, Axiom 6 states that from her current point of view, the decision maker is indifferent between getting, at time $\tau + 1$, a consequence that she cannot describe using her current language, but may be able to describe at time τ, and a consequence that she will still not be able to describe with her time-τ language. Since $\hat{c}$ can also be any consequence discovered between times t and τ, Axiom 6 also postulates that from her current point of view, the decision maker is indifferent between getting, at time $\tau + 1$, different consequences that she cannot currently describe.

The need for Axiom 6 arises from the intertemporal context. It is therefore a conceptually new axiom. The need arises because with a long time-horizon, the decision-maker anticipates that awareness will gradually grow over time. The axiom has a flavor of the decision-maker having rational expectations to the extent possible about her unawareness. By evaluating different currently unknown consequences the same way, the decision maker does not anticipate that she will make systematic adjustments of the utility of currently unknown consequences once they are discovered. In other words, she does not expect to make systematic mistakes in her evaluation of the unknown.[11]

The following axioms connect preferences across different levels of awareness. To state Axiom 7, define, for all $\omega_t \in \mathbf{\Omega}_{T-1}$ and for all $f \in F$,

$$L_{\omega_t}(f) = \{h_{\mathbf{\Omega}_{T-1}(\omega_t)} f | h(\omega_\tau)(s) = l_\tau \in \Delta(C(\omega_t)) \text{ for all } \omega_\tau \in \Omega_\tau(\omega_t),\ s \in S_{\tau+1}(\omega_\tau) \text{ and } \tau \geq t\},$$

and

$$L_{\omega_t}(F) = \bigcup_{f \in F} L_{\omega_t}(f).$$

The objects in $L_{\omega_t}(F)$ return the same lottery, with support being a subset of $\Delta(C(\omega_t))$, in each time $\tau + 1$ state in the continuation path for all $\tau \geq t$, but can return different lotteries at different times. Hence, $L_{\omega_t}(F)$ is a subset of F that involves risk but no subjective uncertainty, and only involves currently known consequences.[12]

Axiom 7 (Time- and Awareness-Invariant Risk Preferences)**.** *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $l \in L_{\omega_t}(F)$, for all $p, p', q, q' \in \Delta(C(\omega_t))$, if for some $\omega_{\hat{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$, and $\tau \geq \hat{t}$ it is true that $p_{\Omega_\tau(\omega_t)} p'_{\Omega_{\tau+1}(\omega_t)} l \succsim_{\omega_{\hat{t}}} q_{\Omega_\tau(\omega_t)} q'_{\Omega_{\tau+1}(\omega_t)} l$, then it is true for every $\omega_{\hat{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$, and $\tau \geq \hat{t}$.*

[11] When the decision-maker indeed does discover new consequences, they will in general be valued differently than expected. The key point is that she does not expect a systematic bias.

[12] In Axiom 7, $p_{\Omega_\tau(\omega_t)} p'_{\Omega_{\tau+1}(\omega_t)} l$ is the intertemporal act obtained from l by replacing, in the continuation path, all restricted Anscombe-Aumann acts originating at time τ by the constant Anscombe-Aumann act p and all those originating at time $\tau + 1$ by the constant Anscombe-Aumann act p'. When $\tau = T - 1$, the notation is interpreted as the act that returns p in all time-T states and agrees with l in all other partial histories.

Axiom 7 requires that the attitude towards known risks is invariant over time and levels of awareness, both for acts that differ in a single period and in two successive periods. It extends the Risk Preference axiom in Epstein and Schneider (2003) to the context of growing awareness.[13] The axiom contains elements that concern preferences within an awareness level as well as elements that link preferences across awareness levels. The part that links preferences across awareness levels is stronger than the Invariant Risk Preferences Axiom from Karni and Vierø (2017), since it also applies for acts that differ across two successive periods. This was beyond the scope of the framework in Karni and Vierø (2017).[14]

Axiom 8 (Invariant Attitude Towards the Unknown)**.** *For all* $f \in F$ *and for all* $\omega_t \in \mathbf{\Omega}_{T-1}$*,*

(i) if $(x_{\omega_t})_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)}f \sim_{\omega_t} (\alpha c^* + (1-\alpha)c_*)_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)}f$ *then* $(x_{\omega_t})_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)}f \sim_{\omega_t} (\alpha c^* + (1-\alpha)c_*)_{S_{\tau+1}(\omega_\tau)\setminus\tilde{S}_{\tau+1}(\omega_\tau)}f$ *for all* $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$*.*

(ii) if $(x_{\omega_t})_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)}f \sim_{\omega_t} (\alpha c^* + (1-\alpha)c_*)_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)}f$ *then* $(x_{(\omega_t,s_{t+1})})_{S_{t+2}(\omega_t,s_{t+1})\setminus\tilde{S}_{t+2}(\omega_t,s_{t+1})}f \sim_{(\omega_t,s_{t+1})} (\alpha c^* + (1-\alpha)c_*)_{S_{t+2}(\omega_t,s_{t+1})\setminus\tilde{S}_{t+2}(\omega_t,s_{t+1})}f$ *for all* $s_{t+1} \in S_{t+1}(\omega_t)$

Axiom 8 states that the decision maker's attitude towards the unknown is invariant to her level of awareness. She does not become more fearful or excited towards the unknown as her awareness evolves. Part (i) states that the decision maker's current attitude towards the unknown is independent of which future history she is considering. Part (ii) states that the attitude towards the unknown remains unchanged as the decision maker's awareness grows.

Axiom 8 is conceptually new and addresses issues that arise from the intertemporal context. Given that there is gradual resolution of uncertainty, the long-horizon model begs the question of how updated preferences evaluate the future after uncertainty has partially resolved. Axiom 8 also considers how different possible futures are evaluated with respect to the unknown. Karni and Vierø (2017) was silent about this.[15]

[13] The difference is that in the present context, the axiom imposes invariance only for risks that exclusively involve known consequences.

[14] Dominiak and Tserenjigmid (2018) show that in Karni and Vierø (2013), the invariant risk preferences Axiom is implied by the other axioms. It is not clear whether this would also be the case with awareness of unawareness, but regardless, in the present context, the axiom is necessary for acts that differ across two successive periods.

[15] The following is a quote from Karni and Vierø (2017): "The main objective of Theorem 2 is the depiction of the evolution of the decision maker's beliefs. To attain this objective it is not necessary to consider the utility of the abstract consequences x_0 and x_1. Therefore, unlike in Theorem 1, in Theorem 2 the domains of the utility functions U_0 and U_1 are $\Delta(C_0)$ and $\Delta(C_1)$, respectively. It is a straightforward exercise to extend the representations in Theorem 2 to include utilities of the abstract consequences x_0 and x_1." The intertemporal model highlights that the exercise is not as straightforward as the quote claims.

Axiom 8 is, perhaps, the most controversial among the axioms. It assumes that the decision maker's attitude towards the unknown, which can be interpreted as her expectations about the average quality of unknown consequences yet to be discovered, is not a function of previous discoveries. The axiom therefore precludes, for example, that drawing a bad consequence from the set of unknown consequences causes the decision maker to consider the remaining unknown consequences as being less bad, or that a draw of a bad consequence is interpreted as a signal that unknown consequences are more likely to be bad. To provide a justification for the axiom, notice that there is a continuum of utility values that each new consequence could potentially be assigned and that they are drawn with replacement in the sense that when a new consequence is assigned a particular utility value, there is still a possibility of drawing a different new consequence in the future with the same utility value. Since the decision maker knows nothing about unknown consequences or the distribution of their utilities, one can argue that she can reasonably maintain her outlook towards the unknown after a discovery and for different future histories. The attitude is simply the part of her preferences that reflects how she feels about discovering new things.

To state the last axiom, the following notation is introduced: For all $\omega_t \in \mathbf{\Omega}_{T-1}$ and for all $s \in S_{t+1}(\omega_t)$, define, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, the event $\mathscr{E}_{\tau+1}(s|\omega_\tau) \subset \Omega_{\tau+1}$ by

$$\begin{aligned}\mathscr{E}_{\tau+1}(s|\omega_\tau) \equiv \big\{(\omega_\tau, s_{\tau+1}) \in \Omega_{\tau+1}(\omega_\tau) : \forall a \in A,\ \text{if } a(s) \in C(\omega_t) \text{ then } a(s_{\tau+1}) \\ = a(s) \text{ and if } a(s) \notin C(\omega_t) \text{ then } a(s_{\tau+1}) \in \{x(C(\omega_\tau))\} \cup (C(\omega_\tau) \setminus C(\omega_t))\big\}.\end{aligned} \tag{14}$$

Definition (14) maps fully describable states into degenerate events and imperfectly describable states into non-degenerate events in $\Omega_{\tau+1}$. The definition can be illustrated using matrices (3) and (4) from Example 2. There, $\mathscr{E}_{t+2}(s_1^1|(\omega_0, s_1^2)) = \{(\omega_0, s_1^2, s_2^1)\}$, $\mathscr{E}_{t+2}(s_1^2|(\omega_0, s_1^2)) = \{(\omega_0, s_1^2, s_2^2), (\omega_0, s_1^2, s_2^5)\}$, $\mathscr{E}_{t+2}(s_1^3|(\omega_0, s_1^2)) = \{(\omega_0, s_1^2, s_2^3), (\omega_0, s_1^2, s_2^7)\}$, and $\mathscr{E}_{t+2}(s_1^4|(\omega_0, s_1^2)) = \{(\omega_0, s_1^2, s_2^4), (\omega_0, s_1^2, s_2^6), (\omega_0, s_1^2, s_2^8), (\omega_0, s_1^2, s_2^9)\}$.

Fix two outcomes $c^*, c_* \in C(\omega_0)$ for which $c^*_{\omega_0} f \succ_{\omega_0} c_{*\omega_0} f$ for some $f \in F$ (hence, given the axioms, for all $f \in F$). Such two outcomes exist by Axiom 4.

Axiom 9 (Forward Awareness Consistency)**.** *For all $f \in F$, for all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $s \in S_{t+1}(\omega_t)$, for all $\omega_{t+1} \in \Omega_{t+1}(\omega_t)$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_{t+1})$, and for all $g, h \in H_{\omega_\tau}(f)$, if*

$$g = (\eta c^* + (1-\eta) c_*)_{S_{\tau+1}(\omega_\tau)} f, \quad \textit{and} \quad h = c^*_{\mathscr{E}_{\tau+1}(s|\omega_\tau)} c_{* S_{\tau+1}(\omega_\tau)} f,$$

then $g \succsim_{\omega_t} h$ if and only if $g \succsim_{\omega_{t+1}} h$.

Axiom 9 assumes that the ranking of objective uncertainty versus subjective uncertainty about future events that are measurable with respect to current awareness is unchanged when moving forward in the event tree. In other words, the ranking of objective versus subjective uncertainty about events the decision maker can currently describe is independent of the level of detail with which the subjective uncertainty can be described. This is imposed for all partial histories following immediately after the current partial history and the states and events in their continuation paths.

Axiom 9 ensures consistency of preferences when looking forward. It is not necessarily reasonable to impose such a requirement looking backwards, since the decision maker's awareness may have reached a higher level. For the same reason, it is only reasonable to impose the requirement for events that are measurable with respect to the decision maker's current set of extended consequences, not for individual states that involve new consequences that the decision maker is currently unaware of. Thus, looking backwards, there are things that the decision maker can take into consideration that she was not able to take into consideration previously. However, looking forward, Axiom 9 requires that preferences will be consistent regarding the currently known and well-understood part of the decision maker's universe.[16] Axiom 9 extends the Awareness Consistency axioms in Karni and Vierø (2017) to the long time-horizon, but it is not conceptually new. Comparing it to Dynamic Consistency in Epstein and Schneider (2003), Axiom 9 is only imposed on a subset of acts, to allow for revisions due to growing awareness.

5 Representation

Theorem 1 provides a representation of preferences over intertemporal acts at each partial history and awareness level. It also connects preferences, through connecting utilities and beliefs, across partial histories and awareness levels. To facilitate reading the theorem, keep the following notation in mind: In the statement of Theorem 1, ω_t is the current partial history at which the preference is expressed, ω_τ is used to denote the partial history in which a restricted Anscombe-Aumann act originates, and s indexes the states in which the uncertainty regarding the restricted Anscombe-Aumann act resolves. To facilitate notation, define, for each partial history ω_τ,

$$\pi_{\omega_t}(\omega_\tau) \equiv \pi_{\omega_t}\left(\{\omega : \mathbf{P}_\tau(\omega) = (\omega_\tau)\}\right).$$

In words, the probability of a partial history equals the probability of the set of full histories that project onto that partial history. The notation $\mathscr{E}_{\tau+1}(s|\omega_t)$ is defined in (14).

[16]There may be situations in which Axiom 9 is too strong. For example, the decision maker could become ambiguity averse in response to increases in awareness. Such a possibility is investigated in Dominiak and Tserenjigmid (2020) for a one-shot increase in awareness and the decision maker being myopic with respect to his unawareness. It is far from clear how ambiguity aversion would interplay with awareness of unawareness or with the long time horizon.

Theorem 1. *For $T \geq 2$, the following statements are equivalent:*

(a) $\{\succsim_{\omega_t}\}_{\omega_t \in \mathbf{\Omega}_{T-1}}$ *satisfy Axioms 1 through 9.*

(b) For all $\omega_t \in \mathbf{\Omega}_{T-1}$, there exist a real-valued, continuous, non-constant Bernoulli-utility function u_{ω_t} on $C(\omega_t)$, a unique probability measure π_{ω_t} on Ω with $\pi_{\omega_t}(\omega) = 0$ if $\omega \notin \Omega(\omega_t)$ and $\pi_{\omega_t}(\omega_{T-1}) > 0$, for all $\omega_{T-1} \in \Omega_{T-1}(\omega_t)$, and a parameter $\beta > 0$ such that for every $\omega_t \in \mathbf{\Omega}_{T-1}$, $\succsim_{\omega_t}$ is represented by $V_{\omega_t}(\cdot)$, where

$$V_{\omega_t}(f) = \sum_{\tau=t}^{T-1} \beta^{\tau-t} \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \sum_{s \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, s) \sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c) u_{\omega_t}(c). \tag{15}$$

The function u_{ω_0} is unique up to positive linear transformations. For all ω_t, the function u_{ω_t} is unique given the transformation of u_{ω_0}, and for all $c \in C(\omega_t)$, $u_{\omega_\tau}(c) = u_{\omega_t}(c)$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$.

The probability measures π_{ω_t} satisfy that for all $\omega_{t+1} \in \Omega_{t+1}(\omega_t)$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_{t+1})$, and for all $s, \tilde{s} \in S_{t+1}(\omega_t)$, we have that

$$\frac{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau))} = \frac{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau))}. \tag{16}$$

Proof: The proof of Theorem 1 is in the appendix.

Remark 1: When T=1, Axioms 7, 8(ii), and 9 are not necessary. They concern how updated preferences evaluate the future after resolution of uncertainty. When T=1, there is no future left to evaluate after the initial resolution of uncertainty. Axioms 1 through 9 are still sufficient for representation by (15) when T=1.

Remark 2: An equivalent representation to that in part (b) of Theorem 1 is that there in addition exists a parameter $u^* \in \Re$, and the utility function in (15) is replaced by

$$V_{\omega_t}(f) = \sum_{\tau=t}^{T-1} \beta^{\tau-t} \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \sum_{s \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, s) \left(\sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c) u_{\omega_t}(c) + \Big(1 - \sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c)\Big) u^* \right). \tag{17}$$

In that case, the Bernoulli utility function u_{ω_0} is unique up to positive affine transformations, not just positive linear transformations. The parameter u^* is unique given the transformation of u_{ω_0} and can be interpreted as the decision-maker's attitude towards the unknown. A high value indicates that she is excited about the unknown, whereas a low value indicates that she is fearful towards the unknown. The normalization $u^* = 0$ in Theorem 1 fixes the level of the Bernoulli utility function since that level now characterizes the attitude towards the unknown.

Note that the sum of probabilities $\sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c)$ can be less than one. It is less than one when the lottery $f(\omega_\tau)(s)$ assigns positive probability to currently unknown consequences. Note also that the updating rule in (16) can equivalently be expressed as

$$\frac{\pi_{\omega_t}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(s|\omega_\tau)\})}{\pi_{\omega_t}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)\})} = \frac{\pi_{\omega_{t+1}}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(s|\omega_\tau)\})}{\pi_{\omega_{t+1}}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)\})}.$$

The representation of preferences over intertemporal acts in (15) has the following form: When finding herself in partial history ω_t, the decision maker acts as if she computes subjective expected utility over future partial histories using her ω_t-beliefs and computes the discounted sum of utilities using the time and history invariant discount factor β. The utility functions u_{ω_t} are time and history, and thus awareness, invariant for consequences that are common to the partial histories. The representation reflects an attitude towards the unknown, which is also time and history, and thus awareness, invariant, since at any partial history, the agent assigns an expected utility of zero to unknown consequenecs. For each state s, the decision maker computes a generalized von Neumann-Morgenstern utility of the lottery that the intertemporal act under evaluation returns in that state. The generalized von Neumann-Morgenstern utility evaluates all outcomes in $C(\omega_t)$ according to u_{ω_t} and collapses all unknown consequences from the ω_t-point of view into one unknown consequence, which is assigned utility value of zero. Act f assigns probability $1-\sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c)$ to this default unknown consequence in state s. The representation (15) is thus more general than standard representations, because the sum of probabilities $\sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c)$ can be less than one.

When awareness grows and new consequences are indeed discovered, the resulting Bernoulli-utility function is an extension of the previous one. The decision maker's attitude towards the unknown remains unchanged at zero in response to the increase in awareness. Hence, she does not become more excited about or fearful towards the unknown. Beliefs are updated according to (16), which requires that relative posterior probabilities of events that are measurable with respect to the prior set of extended consequences agree with the relative prior conditional probabilites of those events. Theorem 1 thus succeeds in separating the evolution of the decision maker's attitude towards the unknown from the evolution of her sense of unawareness. The latter is captured by her beliefs about making new discoveries. The attitude towards the unknown can be interpreted as the decision maker's subjective expected value of yet undiscovered consequences. In Theorem 1 this attitude is normalized to zero.

The invariance of the attitude across awareness levels means that this expectation is unaffected by what is discovered.

Existence, linearity, and state separability of the representation is a result of Axiom 2. That only the continuation path enters (15) follows from Axiom 1. Axiom 3 aides in identifying the subjective probabilities, and the full support of π_{ω_t} on Ω_{T-1} follows from Axiom 4. Axiom 5 ensures that at each history, the attitude towards the unknown is independent of the act under evaluation. Axiom 7 ensures exponential discounting as well as time- and awareness invariance of the discount factor β and that subsequent Bernoulli-utility functions are extensions of preceding ones. The collapsing of all unknown consequences into one, and the time- and awareness invariance of the attitude towards the unknown, are results of Axioms 6 and 8. The updating rule for beliefs follows from (15) and Axiom 9.

The next result in Theorem 2 provides a recursive formulation of utility. However, the decision maker can only forecast her future utility function to the extent of her awareness. That is, she can currently only express her future utility with respect to her current set of extended consequences. She does not yet know what will be her Bernoulli-utility of consequences to be discovered between the current and the next period.

To ease notation, define $U_{\omega_t}(\hat{p}) = \sum_{c \in C(\omega_t)} \hat{p}(c) u_{\omega_t}(c)$. This is the generalized (since the probabilities need not sum to one) von Neumann-Morgenstern utility of the lottery $\hat{p}$.

Theorem 2. *Let $V_{(\omega_t,s)}(f|C(\omega_t))$ be derived from $V_{(\omega_t,s)}(f)$ by setting $u_{(\omega_t,s)}(c) = 0$ for all $c \in C(\omega_t, s) \setminus C(\omega_t)$. Then the representation in part (b) of Theorem 1 implies that*

$$V_{\omega_t}(f) = \sum_{s \in S_{t+1}(\omega_t)} \pi_{\omega_t}(\omega_t, s) \left[U_{\omega_t}(f(\omega_t)(s)) + \beta V_{(\omega_t,s)}(f|C(\omega_t)) \right]. \tag{18}$$

Proof: The proof of Theorem 2 is in the appendix.

The function $V_{(\omega_t,s)}(f|C(\omega_t))$ can be thought of as the decision maker's current estimate of her future utility function, given her current awareness. The estimate treats all consequences that the decision maker will potentially discover between now and the next period as currently unknown consequences. As a result, they are all assigned a utility value of zero.

As awareness (potentially) evolves and we move from one history to the next, beliefs are updated according to (16). This ensures that relative posterior conditional probabilities of events that are measurable with respect to the prior set of extended consequences agree with the relative prior conditional probabilites of those events. This is sufficient for the recursive representation, since the consequences that the decision maker will potentially discover between the current and the next period are "collapsed" into the current unknown consequence. Thus, future lotteries returning different such unknowns with the same probabilities are equivalent from the current point of view. Then the updating of beliefs in (16) implies that next period beliefs agree with current conditional beliefs. Hence, the convenient recursive relation in (18) applies. It generalizes the standard recursive approach to include unawareness. For a comprehensive textbook discussion of the standard recursive approach and some of the models it can be used to analyze, see e.g. Sargent (1987).

6 An intertemporal asset pricing problem

As an example of the relevance of growing awareness for dynamic decision making, consider a representative agent asset pricing problem.[17] There is one consumer, who lives through time $T = 3$ after which the world ends. Hence, $t \in \{0, 1, 2, 3\}$. The consumer owns a productive asset that yields a dividend each period. In partial history ω_t, next period dividends are denoted by the random variable d_{ω_t}. At $t = 0$ the consumer is only aware of one possible dividend realization D_0, but at each point in time, she entertains the possibility that d_{ω_t} may take other values than currently known. However, she has no information about what these may be or whether there are indeed other possible values. Hence, letting $\mathcal{D}_{\omega_t}$ denote the set of known possible dividends in partial history ω_t, the consumer considers the possibility of a dividend $x_{\omega_t} = \neg\mathcal{D}_{\omega_t}$, which is none of the known possible dividends. Let $\widehat{\mathcal{D}}_{\omega_t} = \mathcal{D}_{\omega_t} \cup \{x_{\omega_t}\}$ be the extended set of dividends, which is the support of the random variable d_{ω_t}.

In the context of this model, dividends should be understood abstractly, not literally as a numerical value. There are several possible interpretations of the discovery of a prior unknown value of the dividend. One interpretation is that the asset is a production process where a by-product is discovered to be of use. A different interpretation is that someone finds a new use for the product such that its value changes or that others make a new discovery which makes the product obsolete. Yet another interpretation is that some unexpected and unfamiliar government policy surprise will affect the business or production.

For example, the asset could be producing a particular type of plastic, which in the current market is used to produce water bottles. The known dividend represents business as usual. A positive surprise could be that someone invents a new technology, such as a 3D printer, which allows new uses. A negative surprise could be that someone discovers a toxic component of your product, which makes it unsuitable for water bottle use.

The underlying process governing dividends may or may not be i.i.d. However, since in any case the consumer is not aware of the full support of this process, the dividend is not i.i.d. from the consumer's point of view. Therefore, the random variable denoting the dividend is

[17] This is a finite-horizon Lucas (1978) tree-type asset pricing problem with growing awareness.

Figure 3: History space for the asset pricing problem.

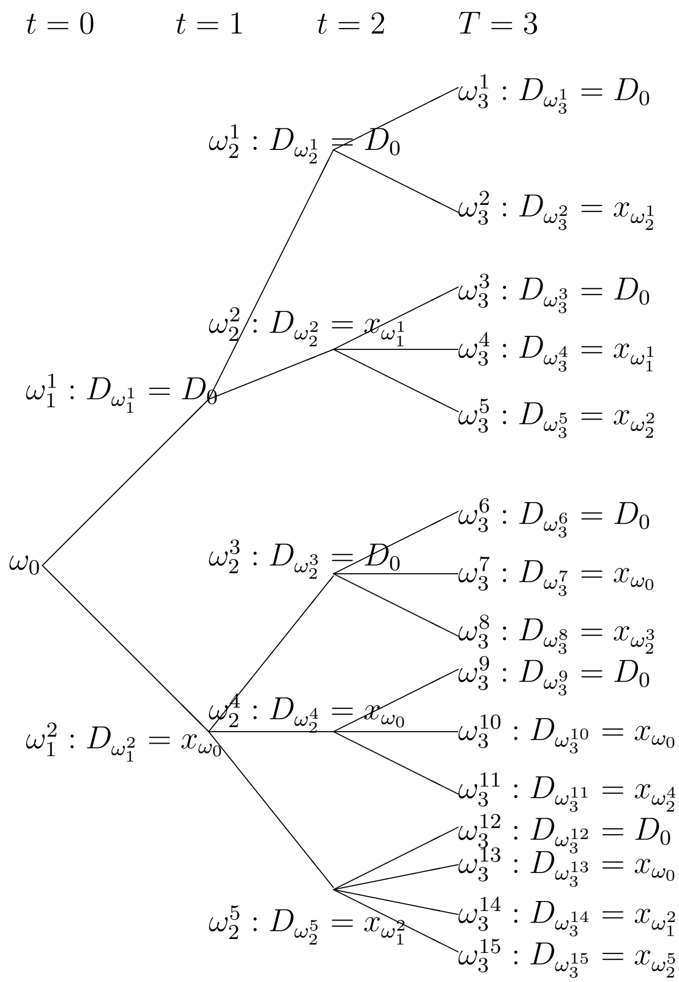

subscripted by the partial history, which also indexes the consumer's level of awareness.

Figure 3 shows the history space for the asset pricing problem. Descriptions of partial histories differ from the ex ante and the ex post perspective. For example, at time 0 the consumer can describe the possibility of a dividend realization at time 1, which is not the known value D_0. However, from the time-0 point of view, the consumer cannot describe this potential new dividend further than this. Hence, it is assigned the value $x_{\omega_0} = \neg\mathcal{D}_{\omega_0}$ from the time-0 perspective. Thus, partial history ω_1^2 can be described, ex ante, as an event in which a new possible dividend realization is discovered at time 1. If such a new dividend is indeed realized at time 1, it takes a concrete value, namely the realization $D_{\omega_1^2}$. This value will replace x_{ω_0} in the tree moving forward from ω_1^2. Following the realization, the set of known dividends is $\mathcal{D}_{\omega_1^2} = \{D_0, D_{\omega_1^2}\}$ and "none of the above" is redefined to $x_{\omega_1^2} = \neg\mathcal{D}_{\omega_1^2}$. Hence, the random variable $d_{\omega_1^2}$, which describes time-2 dividends following ω_1^2, has support $\{D_0, D_{\omega_1^2}, x_{\omega_1^2}\}$ from the ω_1^2-perspective, while its support is $\{D_0, x_{\omega_0}, x_{\omega_1^2}\}$ from the ω_0–

perspective.

Let the asset holdings of the representative consumer at the beginning of period t in partial history ω_t be denoted by $\xi_{\omega_{t-1}}$. Each period unfolds as follows: The consumer observes the dividend realization D_{ω_t}, which may be one of the known, previously realized dividends in $\mathcal{D}_{\omega_{t-1}}$, or a new, previously unknown, dividend. If $D_{\omega_t} \notin \mathcal{D}_{\omega_{t-1}}$, i.e. is a new, previously unknown, dividend, the consumer updates her set of possible dividends to $\mathcal{D}_{\omega_t} = \mathcal{D}_{\omega_{t-1}} \cup \{D_{\omega_t}\}$ and then chooses how to divide her resources between consumption c_{ω_t} and the asset. Hence, $\mathcal{D}_{\omega_t}$ and x_{ω_t} indicate knowledge at the moment period-t decisions are made, i.e. after period-t dividends are realized.

Let the price of consumption be normalized to 1 and let the asset price in partial history ω_t be p_{ω_t}. Given prices, the consumption choice determines the asset holding going into the next period, ξ_{ω_t}, as given by (20) below. The asset price will depend on both the dividend realization and the set of known dividends $\mathcal{D}_{\omega_t}$. The latter dependence is one of the differences from models without growing awareness, and it gives a different price evolution than in models without growing awareness. At a given point in time and for a given dividend realization, the asset price as well as price expectations will be different for different levels of awareness.

The expectation in partial history ω_t of the ω_τ-dividend's payoff (keeping in mind that a period's dividends and discoveries are realized before the period's decisions are made) is

$$E_{\omega_t}[\widehat{D}_{\omega_\tau}] = \begin{cases} \widehat{D}_{\omega_\tau} & \text{if } \widehat{D}_{\omega_\tau} \in \mathcal{D}_{\omega_t} \\ u^*_{\omega_t}(\widehat{D}_{\omega_\tau}) & \text{if } \widehat{D}_{\omega_\tau} \notin \mathcal{D}_{\omega_t}. \end{cases} \tag{19}$$

Having reached partial history ω_t, the consumer anticipates a budget constraint for partial history $\omega_\tau \in \boldsymbol{\Omega}(\omega_t)$, of

$$B_{\omega_t}(\xi^{est(\omega_t)}_{\omega_{\tau-1}}, \omega_\tau) = \{(c^{est(\omega_t)}_{\omega_\tau}, \xi^{est(\omega_t)}_{\omega_\tau}) \mid E_{\omega_t}[p_{\omega_\tau} + \widehat{D}_{\omega_\tau}]\xi^{est(\omega_t)}_{\omega_{\tau-1}} = E_{\omega_t}[p_{\omega_\tau}]\xi^{est(\omega_t)}_{\omega_\tau} + c^{est(\omega_t)}_{\omega_\tau}\}. \tag{20}$$

The notations $c^{est(\omega_t)}_{\omega_\tau}$ and $\xi^{est(\omega_t)}_{\omega_\tau}$ denote the consumer's estimates of, or plans for, respectively, her ω_τ consumption and ω_τ asset holdings when she is at partial history ω_t. The budget constraints involve estimated consumptions and asset holdings, since in partial histories that involve not yet discovered dividends, the consumer can only use her current estimate of what the dividend as well as the asset price will be. Therefore, her choices may eventually differ from planned.

Having reached partial history ω_t, the consumer's problem is to maximize her utility,

$$V_{\omega_t}(\{c^{est(\omega_t)}_{\omega_\tau}\}) = u(c_{\omega_t}) + \sum_{\tau=t+1}^{3} \beta^{\tau-t} \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \pi_{\omega_t}(\omega_\tau) u(c^{est(\omega_t)}_{\omega_\tau}),$$

subject to the budget constraints in (20) for all $\omega_\tau \in \mathbf{\Omega}(\omega_t)$. Like the budget constraints, the objective function involves the consumer's estimate of, or plan for, consumption, rather than actual consumption. Since a new discovery leads to both an immediate income effect and a change in expected future marginal tradeoffs, these estimates will in general change in response to growing awareness.

As usual, the problem can be solved backwards, although with the caveat that estimated values may not be correct, since awareness may evolve. Time 3 is the last period. Hence the time-3 asset price will be zero in any ω_3, and correctly anticipated to be zero in all earlier partial histories. Therefore, consumption in partial history ω_3 will be $c_{\omega_3} = D_{\omega_3}\xi_{\omega_2}$, and earlier estimates of it will be $c_{\omega_3}^{est(\omega_t)} = E_{\omega_t}[\widehat{D}_{\omega_3}\xi_{\omega_2}]$, where the dividend value $\widehat{D}_{\omega_3}$ may change from the subjective estimate to the actual value, and ξ_{ω_2} may change because growing awareness leads to different marginal tradeoffs and, as a result, different choices.

For any partial history $\omega_t \in \mathbf{\Omega_{T-1}}$, the first-order conditions for optimizing in ω_t are

$$u'(c_{\omega_t})p_{\omega_t} = \beta \sum_{\omega_{t+1}\in\Omega_{t+1}(\omega_t)} \pi_{\omega_t}(\omega_{t+1})u'(c_{\omega_{t+1}}^{est(\omega_t)})E_{\omega_t}[\widehat{D}_{\omega_{t+1}} + p_{\omega_{t+1}}], \tag{21}$$

and, for all $\omega_{t+1} \in \Omega_{t+1}(\omega_t)$,

$$\pi_{\omega_t}(\omega_{t+1})u'(c_{\omega_2}^{est(\omega_t)})E_{\omega_t}[p_{\omega_{t+1}}] = \beta \sum_{\omega_{t+2}\in\Omega_{t+2}(\omega_{t+1})} \pi_{\omega_t}(\omega_{t+2})u'(c_{\omega_{t+2}}^{est(\omega_t)})E_{\omega_t}[\widehat{D}_{\omega_{t+2}} + p_{\omega_{t+2}}], \tag{22}$$

with (22) being irrelevant when $t = 2$ (since the world ends after $t = 3$). In ω_0, there are the additional first-order conditions

$$\pi_{\omega_0}(\omega_2)u'(E_{\omega_0}[c_{\omega_2}])E_{\omega_0}[p_{\omega_2}] = \beta \sum_{\omega_3\in\Omega_3(\omega_2)} \pi_{\omega_0}(\omega_3)u'(E_{\omega_0}[c_{\omega_3}])E_{\omega_0}[\widehat{D}_{\omega_3}] \tag{23}$$

for all $\omega_2 \in \Omega_2$.

In a competitive equilibrium, markets are expected to clear and do indeed clear. That is,

$$c_{\omega_t} = D_{\omega_t}, \text{ and } c_{\omega_\tau}^{est(\omega_t)} = E_{\omega_t}[\widehat{D}_{\omega_\tau}] \text{ for all } \omega_t \text{ and } \omega_\tau \in \mathbf{\Omega}(\omega_t), \tag{24}$$

where $E_{\omega_t}[\widehat{D}_{\omega_\tau}]$ is given by (19). The latter part of (24) therefore assumes rational expectations to the extend possible given the level of awareness, using the subjective value of unknowns for partial histories with currently unknown dividend realizations. By the representation result in Theorem 1, all unknowns will be assigned the same subjective value of zero, i.e.

$$u_{\omega_t}^*(\widehat{D}_{\omega_\tau}) = 0 \text{ for all } \omega_t, \omega_\tau, \text{ and } \widehat{D}_{\omega_\tau} \notin \mathcal{D}_{\omega_t}. \tag{25}$$

Also, by the updating rule in (16), we have the following relationships:

$$\frac{\pi_{\omega_0}(\omega_3^1)}{\pi_{\omega_0}(\omega_3^2)} = \frac{\pi_{\omega_1^1}(\omega_3^1)}{\pi_{\omega_1^1}(\omega_3^2)} = \frac{\pi_{\omega_2^1}(\omega_3^1)}{\pi_{\omega_2^1}(\omega_3^2)}, \tag{26}$$

$$\frac{\pi_{\omega_0}(\omega_3^3)}{\pi_{\omega_0}(\omega_3^4)+\pi_{\omega_0}(\omega_3^5)} = \frac{\pi_{\omega_1^1}(\omega_3^3)}{\pi_{\omega_1^1}(\omega_3^4)+\pi_{\omega_1^1}(\omega_3^5)} = \frac{\pi_{\omega_2^2}(\omega_3^3)}{\pi_{\omega_2^2}(\omega_3^4)+\pi_{\omega_2^2}(\omega_3^5)}, \tag{27}$$

$$\frac{\pi_{\omega_0}(\omega_3^6)}{\pi_{\omega_0}(\omega_3^7)+\pi_{\omega_0}(\omega_3^8)} = \frac{\pi_{\omega_1^2}(\omega_3^6)}{\pi_{\omega_1^2}(\omega_3^7)+\pi_{\omega_1^2}(\omega_3^8)}, \tag{28}$$

$$\frac{\pi_{\omega_0}(\omega_3^9)}{\pi_{\omega_0}(\omega_3^{10})+\pi_{\omega_0}(\omega_3^{11})} = \frac{\pi_{\omega_1^2}(\omega_3^9)}{\pi_{\omega_1^2}(\omega_3^{10})+\pi_{\omega_1^2}(\omega_3^{11})}, \text{ and} \tag{29}$$

$$\frac{\pi_{\omega_0}(\omega_3^{12})}{\pi_{\omega_0}(\omega_3^{13})+\pi_{\omega_0}(\omega_3^{14})+\pi_{\omega_0}(\omega_3^{15})} = \frac{\pi_{\omega_1^2}(\omega_3^{12})}{\pi_{\omega_1^2}(\omega_3^{13})+\pi_{\omega_1^2}(\omega_3^{14})+\pi_{\omega_1^2}(\omega_3^{15})}. \tag{30}$$

In addition,

$$\frac{\pi_{\omega_1^2}(\omega_3^6)}{\pi_{\omega_1^2}(\omega_3^7)} = \frac{\pi_{\omega_2^3}(\omega_3^6)}{\pi_{\omega_2^3}(\omega_3^7)} \quad \text{and} \quad \frac{\pi_{\omega_1^2}(\omega_3^6)}{\pi_{\omega_1^2}(\omega_3^8)} = \frac{\pi_{\omega_2^3}(\omega_3^6)}{\pi_{\omega_2^3}(\omega_3^8)}, \tag{31}$$

$$\frac{\pi_{\omega_1^2}(\omega_3^9)}{\pi_{\omega_1^2}(\omega_3^{10})} = \frac{\pi_{\omega_2^4}(\omega_3^9)}{\pi_{\omega_2^4}(\omega_3^{10}))} \quad \text{and} \quad \frac{\pi_{\omega_1^2}(\omega_3^9)}{\pi_{\omega_1^2}(\omega_3^{11})} = \frac{\pi_{\omega_2^4}(\omega_3^9)}{\pi_{\omega_2^4}(\omega_3^{11})}, \tag{32}$$

$$\frac{\pi_{\omega_1^2}(\omega_3^{12})}{\pi_{\omega_1^2}(\omega_3^{13})} = \frac{\pi_{\omega_2^5}(\omega_3^{12})}{\pi_{\omega_2^5}(\omega_3^{13})} \quad \text{and} \quad \frac{\pi_{\omega_1^2}(\omega_3^{12})}{\pi_{\omega_1^2}(\omega_3^{14})+\pi_{\omega_1^2}(\omega_3^{15})} = \frac{\pi_{\omega_2^5}(\omega_3^{12})}{\pi_{\omega_2^5}(\omega_3^{14})+\pi_{\omega_2^5}(\omega_3^{15})}. \tag{33}$$

The first-order conditions (21) through (23) together with the market clearing conditions (24) determine prices and expected prices for each level of awareness. Equations (25) through (33) can then be used to obtain the relationship between prices and expected prices across different levels of awareness. Hence, the evolution of prices and expected prices as time, uncertainty, and awareness unfolds can now be analyzed.

The expected future prices will evolve as awareness increases. Ex-ante, all potential new dividends are considered "none of the above" and are thus assigned the same value, the current subjective estimate of not yet discovered dividends, which is normalized to zero. When awareness expands and a new possible dividend is discovered, the support of future dividends becomes more refined. This newly discovered information will affect future prices as well as expected future prices.

For the following discussion, the consumer's Bernoulli utility is assumed to exhibit constant relative risk aversion. Along the history $(\omega_0, \omega_1^1, \omega_2^1)$ no surprises happen. Therefore, $E_{\omega_0}[p_{\omega_2^1}] = E_{\omega_1^1}[p_{\omega_2^1}] = p_{\omega_2^1}$, i.e. price expectations are constant and correct.

Along the history $(\omega_0, \omega_1^1, \omega_2^2)$ a surprise happens at time 2. If, for example, the realized dividend $D_{\omega_2^2}$ is greater than the subjective estimate of zero, that is, if it is a positive surprise, the asset price p_2^2 will be higher than expected. This is due to current consumption in that partial history being less scarce than estimated, and therefore the relative price of current consumption decreases, which with the normalization means that the asset price increases. Since the surprise occurs at time 2, any earlier price expectations use the subjective estimate $u^* = 0$. Therefore, $E_{\omega_0}[p_{\omega_2^2}] = E_{\omega_1^1}[p_{\omega_2^2}] < p_{\omega_2^2}$, i.e. price expectations are constant, but turn out to be too low.

Along the history $(\omega_0, \omega_1^2, \omega_2^3)$ a surprise happens at time 1, and the prior known dividend of D_0 occurs at time 2. The expectation of the future price p_2^3 will change when the surprise occurs (unless the coefficient of relative risk aversion is 1). The discovery changes the future marginal tradeoff because the support of future dividends now contains the newly discovered dividend $D_{\omega_1^2}$. Hence, the time-2 asset price, and the time-1 expectation thereof, will be different from the prior expectation, even when the asset pays the prior known dividend at time 2. That is, $E_{\omega_0}[p_{\omega_2^3}] \neq E_{\omega_1^2}[p_{\omega_2^3}] = p_{\omega_2^3}$, i.e. price expectations fluctuate in response to a surprise. The direction of the fluctuation depends on the coefficient of relative risk aversion.

Along the history $(\omega_0, \omega_1^2, \omega_2^4)$ a surprise happens at time 1, and the newly discovered dividend is repeated at time 2. Consider, for example, the situation where the realized dividend $D_{\omega_1^2}$ is smaller than the subjective estimate of zero, that is, it is a negative surprise. The expectation of the future price p_2^4 will decrease immediately when the surprise occurs, since ω_2^4-consumption is now scarcer than estimated, causing its relative price to increase, which with the normalization means that the asset price $p_{\omega_2^4}$ decreases. The discovery also changes the future marginal tradeoff because the support of future dividends now contains the newly discovered dividend $D_{\omega_1^2}$, but the income effect dominates. Hence, $E_{\omega_0}[p_{\omega_2^4}] > E_{\omega_1^2}[p_{\omega_2^4}] = p_{\omega_2^4}$, i.e. price expectations again fluctuate. In this particular case, the negative surprise results in a decrease in the expected price. In general, for future partial histories with no further surprises, expectations immediately adjust to the correct level.

Finally, consider the history $(\omega_0, \omega_1^2, \omega_2^5)$. Here, surprises happen at both times 1 and 2. Suppose the first surprise is negative and the second positive. Time-1 expectations about the asset price $E_{\omega_1^2}[p_{\omega_2^5}]$ will turn out to be too low, because the second, positive, surprise means that ω_2^5-consumption will be less scarce than anticipated at time 1. But the expectation of $p_{\omega_2^5}$ also changes in response to the time-1 surprise, that is, $E_1^2[p_2^5] \neq E_0[p_2^5]$, with the direction of change depending on the coefficient of relative risk aversion. The time-1 surprise also affects the time-1 asset price. Unless the coefficient of relative risk aversion is 1, the time-1 asset price in ω_1^2 will be different from its expectation: $p_{\omega_1^2} \neq E_{\omega_0}[p_{\omega_1^2}]$.

Summarizing the discussion above, awareness of unawareness results in the following general pricing rule. The estimated gross return between the current partial history ω_t and next period partial history ω_{t+1} is

$$\widehat{R}_{\omega_{t+1}}(\omega_t) = \frac{E_{\omega_t}[\widehat{D}_{\omega_{t+1}} + p_{\omega_{t+1}}]}{p_{\omega_t}},$$

while the estimated stochastic discount factor is

$$\widehat{\beta}(\omega_t) = \beta \frac{u'(c^{*est(\omega_t)}_{\omega_{t+1}})}{u'(c^*_{\omega_t})},$$

where $c^{*est(\omega_t)}_{\omega_{t+1}}$ and $c^*_{\omega_t}$ are the estimated future, respectively current, equilibrium levels of consumption, and $E_{\omega_t}[\widehat{D}_{\omega_{t+1}}]$ is given by (19). The conditions (21) through (24) therefore imply that

$$E_{\omega_t}[\widehat{\beta}(\omega_t) \times \widehat{R}_{\omega_{t+1}}(\omega_t)] = 1, \tag{34}$$

which is in essence the same risk-bearing rule and resulting pricing rule as in the standard model. The difference between (34) and the standard model is that both the gross return and the stochastic discount factor are affected by awareness of unawareness and growing awareness. However, the relationship can be estimated using the subjective value of the unknown. Pricing then becomes a statistical issue. In response to expansions in awareness, estimates may turn out to be incorrect. As a result, prices and their expectations will adjust over time when the estimated values in (21) through (24) are gradually replaced by their realized values.

In conclusion, prices and their expectations will be more volatile under growing awareness, since surprises will materialize in immediate adjustments of prices and their expectations. Growing awareness reveals new information that cannot be taken into account before it is revealed. Still, the economy is in equilibrium at any node in the event tree, and expectations are, at any node, rational to the extent possible given the limited awareness. In other words, higher price volatility is an equilibrium effect of growing awareness. Finally, it is worth emphasizing that the higher volatility arises in an otherwise standard representative agent model, without the need for heterogeneity.

7 Conclusion

This paper has presented an intertemporal model of growing awareness, which generalizes both the standard event-tree framework and the framework from Karni and Vierø (2017) of awareness of unawareness. In the intertemporal context, unawareness and the decision-maker's cognizance thereof imply that in general the decision-maker will change her course of action in response to awareness growth. Awareness of unawareness also changes her *expectations* about her future course of action in response to awareness growth. The paper has illustrated how expectations evolve as the decision-maker gradually becomes more aware.

At first glance, the problem is seemingly intractable: With a long time horizon, there is a great number of ways in which awareness may grow, both in terms of when increases in awareness occur, what and how much is discovered at any given time, and in which order discoveries are made. The framework provided incorporates all these elements of the problem in a tractable manner.

An axiomatic structure is provided that allows for a representation of preferences over intertemporal acts under awareness of unawareness. The resulting utility function is separable across time and states and has the standard subjective expected utility form as a special case in the absence of awareness of unawareness. With awareness of unawareness present, the decision maker uses a generalized expected utility for each partial history and behaves as if acts were describable with respect to the uncertainties she can express given her current awareness. A recursive formulation of intertemporal utility is also obtained. This recursive formulation makes possible convenient analysis of, and accommodation of awareness and growing awareness in, a large class of problems that use the tools from dynamic programming.

The relevance of growing awareness for dynamic decision making is clear from the intertemporal asset pricing problem considered. When surprises may occur and the agent entertains the possibility that she may encounter surprises, asset prices as well as expectations thereof are more volatile than without growing awareness. Growing awareness reveals new information that cannot be taken into account before it is revealed.

The results in this paper imply that even when facing complex problems with awareness of unawareness and long time horizons, the agent can make complete contingent plans, also for events that involve new discoveries, to the extent that she can describe these plans. The axiomatic structure ensures dynamic consistency in a forward looking way, but not necessarily looking backwards. When awareness does grow, the agent may wish to change her course of action in response to her new awareness. She will, however, still maintain that her original plan was the right one given the awareness she had at the time it was made. Thus, the agent is rational to the extent possible given her limited awareness.

A Proof of Theorem 1

A.1 Sufficiency of Axioms

The set of intertemporal acts F is a convex set, and $\succ_{\omega_t}$ satisfies Axiom 2 for all $\omega_t \in \mathbf{\Omega}_{T-1}$. Thus, by the mixture space theorem, there exists, for all ω_t, a real-valued function $V_{\omega_t} : F \to$

$\Re$ such that $\succ_{\omega_t}$ on F is represented by V_{ω_t} and

$$V_{\omega_t}(\alpha f + (1-\alpha)f') = \alpha V_{\omega_t}(f) + (1-\alpha)V_{\omega_t}(f') \tag{35}$$

for all $f, f' \in F$. Moreover, V_{ω_t} is unique up to positive linear transformation: V'_{ω_t} also represents $\succ_{\omega_t}$ if and only if $V'_{\omega_t} = \kappa V_{\omega_t} + \zeta$, with $\kappa > 0$.

Lemma 1. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, the function V_{ω_t} satisfies*

$$V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} V_{\omega_t}(\omega_\tau)(f(\omega_\tau)),$$

i.e. V_{ω_t} is separable across partial histories.

Remark L1: Note that in Lemma 1, $f(\omega_\tau)$ is the restricted Anscombe-Aumann act that originates in partial history ω_τ.

Proof of Lemma 1: Fix $f^* \in F$ and for each $f \in F$ and $\omega_\tau \in \mathbf{\Omega}_{T-1}$, let $f^{\omega_\tau} = f_{\omega_\tau} f^* \in F$ be defined by $f^{\omega_\tau}(\omega_\tau) = f(\omega_\tau)$ and $f^{\omega_\tau}(\tilde{\omega}_{\tilde{\tau}}) = f^*(\omega)$ for $\tilde{\omega}_{\tilde{\tau}} \neq \omega_\tau$. Let $m \equiv \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} 1$. For any $f \in F$,

$$\frac{1}{m} f + \frac{m-1}{m} f^* = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} \frac{1}{m} f^{\omega_\tau}. \tag{36}$$

By (35) and (36),

$$\frac{1}{m} \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} V_{\omega_t}(f^{\omega_\tau}) = \frac{1}{m} V_{\omega_t}(f) + \frac{m-1}{m} V_{\omega_t}(f^*). \tag{37}$$

For each $\omega_\tau \in \mathbf{\Omega}_{T-1}$, define $V_{\omega_t}(\omega_\tau) : \Delta(C(\omega_\tau))^{\tilde{S}_{\tau+1}(\omega_\tau)} \times \Delta(\widehat{C}(\omega_\tau))^{S_{\tau+1}(\omega_\tau) \setminus \tilde{S}_{\tau+1}(\omega_\tau)} \to \Re$ (this definition embodies the appropriate restriction on the support in the fully describable states) by

$$V_{\omega_t}(\omega_\tau)(g(\omega_\tau)) = V_{\omega_t}(g(\omega_\tau)_{\omega_\tau} f^*) - \frac{m-1}{m} V_{\omega_t}(f^*).$$

For $f \in F$, this definition gives

$$V_{\omega_t}(\omega_\tau)(f(\omega_\tau)) = V_{\omega_t}(f^{\omega_\tau}) - \frac{m-1}{m} V_{\omega_t}(f^*),$$

which implies

$$\frac{1}{m} \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} V_{\omega_t}(\omega_\tau)(f(\omega_\tau)) = \frac{1}{m} \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} V_{\omega_t}(f^{\omega_\tau}) - \frac{m-1}{m} V_{\omega_t}(f^*).$$

Combining with (37) and multiplying by m on both sides, we get

$$V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}} V_{\omega_t}(\omega_\tau)(f(\omega_\tau)).$$

Thus, the representation is additively separable across partial histories.

Lemma 2. *For all $\omega_\tau \notin \mathbf{\Omega}_{T-1}(\omega_t)$, $V_{\omega_t}(\omega_\tau)(f(\omega_\tau)) = k \in \Re$.*

Proof of Lemma 2: This follows from Axiom 1.

Remark L2: Since k will cancel out when comparing acts, one can set $k = 0$ without affecting anything. For ease of notation, this is adopted.

Lemma 3. *For all* $\omega_t \in \mathbf{\Omega}_{T-1}$,

$$V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(f(\omega_\tau)), \tag{38}$$

with $\rho_{\omega_t}(\omega_\tau) > 0$ *for all* $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$.

Remark L3: In Lemma 3, $f(\omega_\tau)$ is, again, the restricted Anscombe-Aumann act that originates in partial history ω_τ. Therefore, an equivalent way to state (38) is as $V_{\omega_t}(\omega_\tau)(f(\omega_\tau)) = \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(f(\omega_\tau))$.

Proof of Lemma 3: This follows from the Anscombe and Aumann Theorem and Axioms 1, 2, 3(iii), and 4, as will now be shown. By Lemmas 1 and 2,

$$V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} V_{\omega_t}(\omega_\tau)(f(\omega_\tau)). \tag{39}$$

Consider the set of acts whose lottery supports are restricted to $C(\omega_t)$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. By Axiom 4 there is a nonnull one-step-ahead resolution of uncertainty for all partial histories. By Axiom 3(iii) and (39),[18]

$$\begin{aligned} V_{\omega_t}(\omega_\tau)(p) > V_{\omega_t}(\omega_\tau)(q) \Leftrightarrow & \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} V_{\omega_t}(\omega_\tau)(p) > \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} V_{\omega_t}(\omega_\tau)(q) \\ \Leftrightarrow & V_{\omega_t}(\omega'_{\tau'})(p) > V_{\omega_t}(\omega'_{\tau'})(q) \end{aligned} \tag{40}$$

for all $\omega_\tau, \omega'_{\tau'} \in \mathbf{\Omega}_{T-1}(\omega_t)$. Thus, $V_{\omega_t}(\omega_\tau)$ and $V_{\omega_t}(\omega'_{\tau'})$ are ordinally equivalent when evaluating constant restricted Anscombe-Aumann acts whose lottery supports are confined to $C(\omega_t)$ for all $\omega_\tau, \omega'_{\tau'} \in \mathbf{\Omega}_{T-1}(\omega_t)$.

Let $v_{\omega_t} \equiv V_{\omega_t}(\omega_t)$. Then, by ordinal equivalence, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, $V_{\omega_t}(\omega_\tau) = \kappa_{\omega_\tau} v_{\omega_t} + \eta_{\omega_\tau}$, with $\kappa_{\omega_\tau}, \eta_{\omega_\tau} \in \Re$ and $\kappa_{\omega_\tau} > 0$ when restricted to such acts. Hence, by (39), for f, g with $f(\omega_\tau(s)) = p \in \Delta(C(\omega_t))$ and $g(\omega_\tau(s)) = q \in \Delta(C(\omega_t))$ for all $s \in S_{\tau+1}(\omega_\tau)$,

$$f \succ_{\omega_t} g \Leftrightarrow \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \kappa_{\omega_\tau} v_{\omega_t}(f(\omega_\tau)) + \eta_{\omega_\tau} > \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \kappa_{\omega_\tau} v_{\omega_t}(g(\omega_\tau)) + \eta_{\omega_\tau}.$$

[18]Here, the notation p is abused to denote the restricted Anscombe-Aumann act for which $f(\omega_\tau) = p$ for all $s \in S_{\tau+1}(\omega_\tau)$.

Cancel out terms, divide both sides by $\sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \kappa_{\omega_\tau}$, and define $\rho_{\omega_t}(\omega_\tau) \equiv \frac{\kappa_{\omega_\tau}}{\sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \kappa_{\omega_\tau}}$. Then

$$V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(f(\omega_\tau)). \tag{41}$$

By Axiom 4, $\rho_{\omega_t}(\omega_\tau) > 0$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$.

For general acts, it follows from (41) that $V_{\omega_t}(f) = \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(f(\omega_\tau))$ and that $v_{\omega_t}(\omega_\tau)$ and $v_{\omega_t}(\omega'_{\tau'})$ agree when evaluating acts whose lottery supports are restricted to $C(\omega_t)$.

Lemma 4. *For all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$,*

$$\begin{aligned} v_{\omega_t}(\omega_\tau)(f(\omega_\tau)) = & \sum_{s \in \widetilde{S}_{\tau+1}(\omega_\tau)} \mu_{\omega_t}(s|\omega_\tau) \sum_{c \in C(\omega_\tau)} f(\omega_\tau)(s)(c) u_{\omega_t}(\omega_\tau)(c) \\ & + \sum_{s \in S_{\tau+1}(\omega_\tau) \setminus \widetilde{S}_{\tau+1}(\omega_\tau)} \mu_{\omega_t}(s|\omega_\tau) \sum_{\hat{c} \in \widehat{C}(\omega_\tau)} f(\omega_\tau)(s)(\hat{c}) u^*_{\omega_t}(\omega_\tau)(\hat{c}) \end{aligned} \tag{42}$$

where u_{ω_t} and $u^*_{\omega_t}$ are unique up to positive linear transformations and agree on $C(\omega_\tau)$.

Remark L4: Since u_{ω_t} and $u^*_{\omega_t}$ agree on $C(\omega_\tau)$,

$$\begin{aligned} & \sum_{\hat{c} \in \widehat{C}(\omega_\tau)} f(\omega_\tau)(s)(\hat{c}) u^*_{\omega_t}(\omega_\tau)(\hat{c}) \\ & = \sum_{c \in C(\omega_\tau)} f(\omega_\tau)(s)(c) u_{\omega_t}(\omega_\tau)(c) + f(\omega_\tau)(s)(x(C(\omega_\tau))) u^*_{\omega_t}(\omega_\tau)(x(C(\omega_\tau))). \end{aligned}$$

Proof of Lemma 4: First note that Axioms 2, 3(i), 3(ii), 4, and 5 all hold on $H_{\omega_\tau}(f)$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ and all $f \in F$.

Consider $h, h' \in H_{\omega_\tau}(f)$. By Lemma 1, the terms in the utilities of h and h' cancel out for all restricted Anscombe-Aumann acts except the ones originating at ω_τ, since h and h' agree except for that restricted Anscombe-Aumann act. Thus, the choice of conditioning act f is immaterial and, by Lemma 3,

$$h \succ_{\omega_t} h' \Leftrightarrow v_{\omega_t}(\omega_\tau)(h(\omega_\tau)) > v_{\omega_t}(\omega_\tau)(h'(\omega_\tau)).$$

Since $F(\omega_\tau)$ is a convex set, arguments analogous to those preceding Lemma 1 and in the proof of Lemma 1 imply that

$$v_{\omega_t}(\omega_\tau)(h(\omega_\tau)) = \sum_{s \in S_{\tau+1}(\omega_\tau)} v_{\omega_t}(\omega_\tau)(s)(h(\omega_\tau)(s)).$$

The standard induction argument shows that for $p \in \Delta(C(\omega_\tau))$ and $s \in S_{\tau+1}(\omega_\tau)$,

$$v_{\omega_t}(\omega_\tau)(s)(p) = \sum_{c \in C(\omega_\tau)} p(c) u_{\omega_t}(\omega_\tau)(s)(c),$$

with $u_{\omega_t}(\omega_\tau)(s)(c) = v_{\omega_t}(\omega_\tau)(s)(c)$, where the former c denotes the consequence c and the latter c denotes the lottery that returns c with probability 1.

Similar arguments show that for $s \in S_{\tau+1}(\omega_\tau) \setminus \widetilde{S}_{\tau+1}(\omega_\tau)$ and $\hat{p} \in \Delta(\widehat{C}(\omega_\tau))$,

$$v_{\omega_t}(\omega_\tau)(s)(\hat{p}) = \sum_{\hat{c} \in \widehat{C}(\omega_\tau)} \hat{p}(\hat{c}) u^*_{\omega_t}(\omega_\tau)(s)(\hat{c}),$$

where $u^*_{\omega_t}(\omega_\tau)(s)(\hat{c}) = v_{\omega_t}(\omega_\tau)(s)(\hat{c})$.

Let $\mathbf{H}_{\omega_\tau}(f) \equiv \{h_{\omega_\tau} f | h : S_{\tau+1}(\omega_\tau) \to \Delta(C(\omega_\tau))\}$, i.e. the subset of $H_{\omega_\tau}(f)$ for which the support of the lotteries in h are restricted to $C(\omega_\tau)$. Consider $h, h' \in \mathbf{H}_{\omega_\tau}(f)$. By Lemma 1, the choice of conditioning act f is immaterial. By Axiom 4, there exists at least one $\succ_{\omega_t}$-nonnull state $s' \in S_{\tau+1}(\omega_\tau)$. By Axiom 3(i), for any $p, q \in \Delta(C(\omega_\tau))$,

$$\begin{aligned} &\sum_{c \in C(\omega_\tau)} p(c) u_{\omega_t}(\omega_\tau)(s)(c) > \sum_{c \in C(\omega_\tau)} q(c) u_{\omega_t}(\omega_\tau)(s)(c) \\ \Leftrightarrow &\sum_{c \in C(\omega_\tau)} p(c) u_{\omega_t}(\omega_\tau)(s')(c) > \sum_{c \in C(\omega_\tau)} q(c) u_{\omega_t}(\omega_\tau)(s')(c) \end{aligned}$$

for all $\succ_{\omega_t}$-nonnull $s \in S_{\tau+1}(\omega_\tau)$. Thus, standard arguments following those in the proof of Lemma 3 imply that there exists a unique probability measure $\mu_{\omega_t}(\cdot|\omega_\tau)$ on $S_{\tau+1}(\omega_\tau)$ such that for $h, h' \in \mathbf{H}_{\omega_\tau}(f)$

$$\begin{aligned} h \succ_{\omega_t} h' \Leftrightarrow &\sum_{s \in S_{\tau+1}(\omega_\tau)} \mu_{\omega_t}(s|\omega_\tau) \sum_{c \in C(\omega_\tau)} h(\omega_\tau)(s)(c) u_{\omega_t}(\omega_\tau)(c) \\ > &\sum_{s \in S_{\tau+1}(\omega_\tau)} \mu_{\omega_t}(s|\omega_\tau) \sum_{c \in C(\omega_\tau)} h'(\omega_\tau)(s)(c) u_{\omega_t}(\omega_\tau)(c), \end{aligned}$$

recalling that by Lemma 1 the choice of conditioning act f is immaterial.

Analogous arguments to those above (using Axiom 3(ii) in place of 3(i)) imply that there exists a unique probability measure $\phi_{\omega_t}(\cdot|\omega_\tau)$ on $S_{\tau+1}(\omega_\tau) \setminus \tilde{S}_{\tau+1}(\omega_\tau)$ such that for all $h, h' \in H_{\omega_t}(F)$ that agree in all $s \in \tilde{S}_{\tau+1}(\omega_\tau)$,

$$\begin{aligned} h \succ_{\omega_t} h' \Leftrightarrow &\sum_{s \in S_{\tau+1}(\omega_\tau) \setminus \tilde{S}_{\tau+1}(\omega_\tau)} \phi_{\omega_t}(s|\omega_\tau) \sum_{\hat{c} \in \hat{C}(\omega_\tau)} h(\omega_\tau)(s)(\hat{c}) u^*_{\omega_t}(\omega_\tau)(\hat{c}) \\ > &\sum_{s \in S_{\tau+1}(\omega_\tau) \setminus \tilde{S}_{\tau+1}(\omega_\tau)} \phi_{\omega_t}(s|\omega_\tau) \sum_{\hat{c} \in \hat{C}(\omega_\tau)} h'(\omega_\tau)(s)(\hat{c}) u^*_{\omega_t}(\omega_\tau)(\hat{c}). \end{aligned}$$

Now, arguments analogous to those in the proof of Theorem 1 in Karni and Vierø (2017) complete the proof of Lemma 4.

Lemma 5. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$ and all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, $u_{\omega_t}(\omega_\tau)(c) = u_{\omega_t}(\omega_t)(c) \equiv u_{\omega_t}(c)$ for all $c \in C(\omega_t)$.*

Proof of Lemma 5: By the arguments preceeding (41), the functions $v_{\omega_t}(\omega_t)(\cdot)$ and $v_{\omega_t}(\omega_\tau)(\cdot)$ are ordinally equivalent when evaluating constant restricted Anscombe-Aumann acts with support in $C(\omega_t)$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. Hence, $u_{\omega_t}(\omega_t)(\cdot)$ and $u_{\omega_t}(\omega_\tau)(\cdot)$ in Lemma 4 must be equal on $C(\omega_t)$ after suitable linear transformations.

Lemma 6. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$ and all $\hat{\omega}_{\hat{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$, $u_{\hat{\omega}_{\hat{t}}}(c) = u_{\omega_t}(c)$ for all $c \in C(\omega_t)$.*

Proof of Lemma 6: By Lemma 5, $u_{\omega_t}(\omega_\tau)(c) \equiv u_{\omega_t}(c)$ for all $c \in C(\omega_t)$ and all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. Hence, it suffices to consider lottery acts that only differ in the restricted Anscombe-Aumann acts that originate in period $t+1$. Consider $l \in L_{\omega_t}(F)$ and $p, q, p' \in \Delta(C(\omega_t))$. By Axiom 7, it holds that for all $\hat{\omega}_{\hat{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$,

$$\begin{aligned} & p_{\Omega_{t+1}(\omega_t)} p'_{\Omega_{t+2}(\omega_t)} l \succsim_{\omega_t} q_{\Omega_{t+1}(\omega_t)} p'_{\Omega_{t+2}(\omega_t)} l \\ \Leftrightarrow & p_{\Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}})} p'_{\Omega_{\hat{t}+2}(\hat{\omega}_{\hat{t}})} l \succsim_{\hat{\omega}_{\hat{t}}} q_{\Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}+1})} p'_{\Omega_{\hat{t}+2}(\omega_{\hat{t}})} l. \end{aligned}$$

Hence, since terms cancel out for all other times than t and $\hat{t}$ respectively,

$$\begin{aligned} & \sum_{\omega_{t+1} \in \Omega_{t+1}(\omega_t)} \rho_{\omega_t}(\omega_{t+1}) v_{\omega_t}(\omega_{t+1})(p) \geq \sum_{\omega_{t+1} \in \Omega_{t+1}(\omega_t)} \rho_{\omega_t}(\omega_{t+1}) v_{\omega_t}(\omega_{t+1})(q) \\ \Leftrightarrow & \sum_{\hat{\omega}_{\hat{t}+1} \in \Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}})} \rho_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1}) v_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1})(p) \geq \sum_{\hat{\omega}_{\hat{t}+1} \in \Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}})} \rho_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1}) v_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1})(q). \end{aligned} \tag{43}$$

By Lemma 5, $v_{\omega_t}(\omega_{t+1})(\cdot) = v_{\omega_t}(\cdot)$ for all ω_{t+1} and $v_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1})(\cdot) = v_{\hat{\omega}_{\hat{t}}}(\cdot)$ for all $\hat{\omega}_{\hat{t}+1}$ when evaluating lotteries with support in $\Delta(C(\omega_t))$. Hence, (43) is equivalent to

$$\begin{aligned} & v_{\omega_t}(p) \sum_{\omega_{t+1} \in \Omega_{t+1}(\omega_t)} \rho_{\omega_t}(\omega_{t+1}) \geq v_{\omega_t}(q) \sum_{\omega_{t+1} \in \Omega_{t+1}(\omega_t)} \rho_{\omega_t}(\omega_{t+1}) v_{\omega_t}(p) \\ \Leftrightarrow & v_{\hat{\omega}_{\hat{t}}}(p) \sum_{\hat{\omega}_{\hat{t}+1} \in \Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}})} \rho_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1}) \geq v_{\hat{\omega}_{\hat{t}}}(q) \sum_{\hat{\omega}_{\hat{t}+1} \in \Omega_{\hat{t}+1}(\hat{\omega}_{\hat{t}})} \rho_{\hat{\omega}_{\hat{t}}}(\hat{\omega}_{\hat{t}+1}). \end{aligned} \tag{44}$$

The expression in (44) is equivalent to

$$v_{\omega_t}(p) \geq v_{\omega_t}(q) \Leftrightarrow v_{\hat{\omega}_{\hat{t}}}(p) \geq v_{\hat{\omega}_{\hat{t}}}(q).$$

Thus, v_{ω_t} and $v_{\hat{\omega}_{\hat{t}}}$ are ordinally equivalent for $p \in \Delta(C(\omega_t))$ for all $\hat{\omega}_{\hat{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$. Hence, after suitable linear transformation, $u_{\hat{\omega}_{\hat{t}}}(c)$ and $u_{\omega_t}(c)$ must be equal on $\Delta(C(\omega_t))$.

Lemma 7. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$ and all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, $\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) = (\beta_{\omega_t})^{\tau - t} \rho_{\omega_t}(\omega_t)$ for some $\beta_{\omega_t} > 0$.*

Proof of Lemma 7: By Axiom 7, for $p, q, p', q' \in \Delta(C(\omega_t))$, and $\hat{\tau} \geq t$

$$p_{\Omega_\tau} p'_{\Omega_{\tau+1}} l \succsim_{\omega_t} q_{\Omega_\tau} q'_{\Omega_{\tau+1}} l \succsim_{\omega_t} \Leftrightarrow p_{\Omega_{\hat{\tau}}} p'_{\Omega_{\hat{\tau}+1}} l \succsim_{\omega_t} q_{\Omega_{\hat{\tau}}} q'_{\Omega_{\hat{\tau}+1}} l, \tag{45}$$

which is equivalent to

$$\begin{aligned}
&\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(p) + \sum_{\omega_{\tau+1} \in \Omega_{\tau+1}(\omega_t)} \rho_{\omega_t}(\omega_{\tau+1}) v_{\omega_t}(\omega_{\tau+1})(p') \\
\geq &\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(q) + \sum_{\omega_{\tau+1} \in \Omega_{\tau+1}(\omega_t)} \rho_{\omega_t}(\omega_{\tau+1}) v_{\omega_t}(\omega_{\tau+1})(q') \\
\Leftrightarrow &\sum_{\omega_{\hat{\tau}} \in \Omega_{\hat{\tau}}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}}) v_{\omega_t}(\omega_{\hat{\tau}})(p) + \sum_{\omega_{\hat{\tau}+1} \in \Omega_{\hat{\tau}+1}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}+1}) v_{\omega_t}(\omega_{\hat{\tau}+1})(p') \\
\geq &\sum_{\omega_{\hat{\tau}} \in \Omega_{\hat{\tau}}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}}) v_{\omega_t}(\omega_{\hat{\tau}})(q) + \sum_{\omega_{\hat{\tau}+1} \in \Omega_{\hat{\tau}+1}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}+1}) v_{\omega_t}(\omega_{\hat{\tau}+1})(q').
\end{aligned} \tag{46}$$

By Lemma 5, $v_{\omega_t}(\omega_\tau)(\cdot) = v_{\omega_t}(\cdot)$ when evaluating lotteries in $\Delta(C(\omega_t))$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. Hence, (46) can be written as

$$\begin{aligned}
&v_{\omega_t}(p) \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) + v_{\omega_t}(p') \sum_{\omega_{\tau+1} \in \Omega_{\tau+1}(\omega_t)} \rho_{\omega_t}(\omega_{\tau+1}) \\
\geq &v_{\omega_t}(q) \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) + v_{\omega_t}(q') \sum_{\omega_{\tau+1} \in \Omega_{\tau+1}(\omega_t)} \rho_{\omega_t}(\omega_{\tau+1}) \\
\Leftrightarrow &v_{\omega_t}(p) \sum_{\omega_{\hat{\tau}} \in \Omega_{\hat{\tau}}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}}) + v_{\omega_t}(p') \sum_{\omega_{\hat{\tau}+1} \in \Omega_{\hat{\tau}+1}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}+1}) \\
\geq &v_{\omega_t}(q) \sum_{\omega_{\hat{\tau}} \in \Omega_{\hat{\tau}}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}}) + v_{\omega_t}(q') \sum_{\omega_{\hat{\tau}+1} \in \Omega_{\hat{\tau}+1}(\omega_t)} \rho_{\omega_t}(\omega_{\hat{\tau}+1}).
\end{aligned} \tag{47}$$

For ease of notation, for all ω_t and all $\tau \geq t$, let

$$\rho_{\omega_t}(\tau) \equiv \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau)$$

Define

$$W_{\omega_t}(\tau)(p, p') \equiv v_{\omega_t}(p)\rho_{\omega_t}(\tau) + v_{\omega_t}(p')\rho_{\omega_t}(\tau + 1).$$

Then (47) implies that $W_{\omega_t}(\tau)$ and $W_{\omega_t}(\hat{\tau})$ are ordinally equivalent for all $\tau, \hat{\tau} \geq t$. By ordinal equivalence of $W_{\omega_t}(\tau)$ and $W_{\omega_t}(\hat{\tau})$,

$$W_{\omega_t}(\hat{\tau}) = \beta_{\omega_t,\tau,\hat{\tau}} W_{\omega_t}(\tau) + \gamma \tag{48}$$

for some $\beta_{\omega_t,\tau,\hat{\tau}} > 0$ and $\gamma \in \Re$. Specifically, (48) holds when $\hat{\tau} = \tau + 1$. Thus, letting $\beta_{\omega_t,\tau} \equiv \beta_{\omega_t,\tau,\tau+1}$,

$$v_{\omega_t}(p)\rho_{\omega_t}(\tau + 1) + v_{\omega_t}(p')\rho_{\omega_t}(\tau + 2) = \beta_{\omega_t,\tau} \left[v_{\omega_t}(p)\rho_{\omega_t}(\tau) + v_{\omega_t}(p')\rho_{\omega_t}(\tau + 1)\right] + \gamma$$

It follows that $\gamma = 0$ and $\rho_{\omega_t}(\tau+1) = \beta_{\omega_t}\rho_{\omega_t}(\tau)$ and $\rho_{\omega_t}(\tau+2) = (\beta_{\omega_t})^2 \rho_{\omega_t}(\tau)$. Then, since τ was chosen arbitrarily, it follows that

$$\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \rho_{\omega_t}(\omega_\tau) = (\beta_{\omega_t})^{\tau - t} \rho_{\omega_t}(\omega_t).$$

Define

$$\tilde{\rho}_{\omega_t}(\omega_\tau) \equiv \frac{\rho_{\omega_t}(\omega_\tau)}{\rho_{\omega_t}(\omega_t)}.$$

Then $\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \tilde{\rho}_{\omega_t}(\omega_\tau) = (\beta_{\omega_t})^{\tau - t}$.

Lemma 8. *For all* $\omega_t \in \mathbf{\Omega}_{T-1}$, $\beta_{\omega_t} = \beta > 0$.

Proof of Lemma 8: For all $p, q, p', q' \in \Delta(C(\omega_t))$ and for all $\hat{\omega}_{\hat{t}}, \tilde{\omega}_{\tilde{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$, $\tau \geq \hat{t}$, and $\tau' \geq \tilde{t}$, it holds, by Axiom 7, that

$$p_{\Omega_\tau} p'_{\Omega_{\tau+1}} l \succsim_{\hat{\omega}_{\hat{t}}} q_{\Omega_\tau} q'_{\Omega_{\tau+1}} l \Leftrightarrow p_{\Omega_{\tau'}} p'_{\Omega_{\tau'+1}} l \succsim_{\tilde{\omega}_{\tilde{t}}} q_{\Omega_{\tau'}} q'_{\Omega_{\tau'+1}} l. \tag{49}$$

By Lemmas 3, 6, and 7, (49) implies that

$$\begin{aligned} &(\beta_{\hat{\omega}_{\hat{t}}})^{\tau - \hat{t}} v(p) + (\beta_{\hat{\omega}_{\hat{t}}})^{\tau + 1 - \hat{t}} v(p') \geq (\beta_{\hat{\omega}_{\hat{t}}})^{\tau - \hat{t}} v(q) + (\beta_{\hat{\omega}_{\hat{t}}})^{\tau + 1 - \hat{t}} v(q') \\ \Leftrightarrow &(\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau' - \tilde{t}} v(p) + (\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau' + 1 - \tilde{t}} v(p') \geq (\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau' - \tilde{t}} v(q) + (\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau' + 1 - \tilde{t}} v(q'). \end{aligned} \tag{50}$$

Consider $\tau, \hat{t}, \tau', \tilde{t}$ such that $\tau - \hat{t} = \tau' - \tilde{t}$. Then (50) implies that

$$(\beta_{\hat{\omega}_{\hat{t}}})^{\tau - \hat{t}} v(p) + (\beta_{\hat{\omega}_{\hat{t}}})^{\tau + 1 - \hat{t}} v(p') = (\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau - \hat{t}} v(p) + (\beta_{\tilde{\omega}_{\tilde{t}}})^{\tau + 1 - \hat{t}} v(p'),$$

which implies that $\beta_{\hat{\omega}_{\hat{t}}} = \beta_{\tilde{\omega}_{\tilde{t}}} \equiv \beta$.

Remark and a definition: Define

$$\mu_{\omega_t}(\omega_\tau) \equiv \frac{\rho_{\omega_t}(\omega_\tau)}{(\beta_{\omega_t})^{\tau - t} \rho_{\omega_t}(\omega_t)}. \tag{51}$$

Notice that $\mu_{\omega_t}(\omega_\tau) \geq 0$ for all ω_τ and that $\sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \mu_{\omega_t}(\omega_\tau) = 1$. Hence, $\mu_{\omega_t}(\omega_\tau)$ is a probability distribution over the time-τ partial histories in the continuation path. Using the definition in (51), the preceeding Lemmas imply that

$$V_{\omega_t}(f) = (\beta_{\omega_t})^{\tau - t} \rho_{\omega_t}(\omega_t) \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \mu_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(f(\omega_\tau))$$

Since V_{ω_t} is unique up to positive linear transformations, we can divide by $\rho_{\omega_t}(\omega_t)$. For notational simplicity, the resulting utility function is also denoted V_{ω_t}. Thus, it has been established from the preceding Lemmas that

$$V_{\omega_t}(f) = (\beta_{\omega_t})^{\tau - t} \sum_{\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)} \mu_{\omega_t}(\omega_\tau) v_{\omega_t}(\omega_\tau)(f(\omega_\tau)).$$

The following Lemma establishes that the probabilities $\mu_{\omega_t}(\omega_\tau)$ defined in (51) are consistent along each history. For $\omega = (\omega_{T-1}, s)$, let $\pi_{\omega_t}(\omega) = \mu_{\omega_t}(\omega_{T-1})\mu_{\omega_t}(s|\omega_{T-1})$, where $\mu_{\omega_t}(\omega_{T-1})$ is defined in (51) and $\mu_{\omega_t}(s|\omega_{T-1})$ is defined in Lemma 4.

Lemma 9. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, $\mu_{\omega_t}(\omega_\tau) = \pi_{\omega_t}(\{\omega|\mathbf{P}_\tau(\omega) = \omega_\tau\})$.*

Proof of Lemma 9: Define $U_{\omega_t}(p) = \sum_{c \in C(\omega_t)} p(c) u_{\omega_t}(c)$. I.e. $U_{\omega_t}(p)$ is the von Neumann-Morgenstern utility of p.

Let F^0 be the set of intertemporal acts for which the support of the lotteries is restricted to $C(\omega_0)$ in any partial history.

Define $X = U_{\omega_0}(\Delta(C_0)))$, that is, X is the set of possible von Neumann-Morgenstern utilities of time-0 consequence lotteries evaluated with the time-0 utility function.

Consider the domain $D = \{f^T : \omega \to \Delta(X)\}$. I.e. D is the set of functions from full histories to the set of lotteries over von Neumann-Morgenstern utilites. For generic element Ψ, denote by $E\Psi(\omega)$ the mean of the lottery $\Psi(\omega)$.

Define

$$\lambda_\tau = \frac{\beta^{\tau-1}}{1 + \beta + \beta^2 + \cdots + \beta^{T-1}}$$

Note that $\sum_{\tau=1}^{T} \lambda_\tau = 1$.

Define $\mathbf{W} : D \to \Re$ by

$$\mathbf{W}(\Psi) = V_{\omega_0}(f)$$

for any $f \in F^0$ satisfying that for any ω,

$$E\Psi(\omega) = U_{\omega_0}\left(\sum_{\tau=1}^{T} \lambda_\tau f(\mathbf{P}_{\tau-1}(\omega))(\mathbf{P}_\tau(\omega))\right). \tag{52}$$

The sum in (52) is a convex combination of the lotteries f returns along history ω.

The function $\mathbf{W}$ is well-defined by Axiom 4 and the properties of V_{ω_0} carry over such that $\mathbf{W}$ admits an expected utility representation:

$$\mathbf{W}(\Psi) = \sum_{\omega \in \Omega} Q(\omega)\Psi(\omega)$$

for some probability measure Q.

Let p^* be a lottery for which $U_{\omega_0}(p^*) = 0$ and let $f^* = p^*$ for all $\tau \neq t$. Consider now acts in $H_{\Omega_t}(f^*) \cap F^0$. For $f \in H_{\Omega_t}(f^*)$,

$$V_{\omega_0}(f) = \beta^t \sum_{\omega_t \in \Omega_t} \mu_{\omega_0}(\omega_t) \sum_{s \in S_{t+1}(\omega_t)} \pi_{\omega_0}(s|\omega_t) U_{\omega_0}(f(\omega_t)(s)).$$

Consider acts in $H_{\Omega_t}(F^*) \cap F^0$ for which $f(\omega) = f(\omega_t)$ for all ω for which $\mathbf{P}_t(\omega) = \omega_t$. These are acts that are measurable w.r.t. the time-t filtration.

Consider a such measurable act $g \in H_{\Omega_t}(F^*) \cap F^0$ for which $g(\omega_t)(s) = \sum_{\tau=1}^{T} \lambda_t f(\mathbf{P}_{\tau-1}(\omega))(\mathbf{P}_\tau(\omega))$ where f satisfies (52). Then

$$V_{\omega_t}(g) = \beta^t \sum_{\omega_t \in \Omega_t} \mu_{\omega_0}(\omega_t) \sum_{s \in S_{t+1}(\omega_t)} \pi_{\omega_0}(s|\omega_t) U_{\omega_0}\left(\sum_{\tau=1}^{T} \lambda_t f(\mathbf{P}_{\tau-1}(\omega))(\mathbf{P}_\tau(\omega))\right).$$

However, we also have that $V_{\omega_0}(g) = \mathbf{W}(\Psi)$, given definition (52), using that linearity of U_{ω_0} implies that $U_{\omega_0}\left(\sum_{\tau=1}^{T} \lambda_t f(\mathbf{P}_{\tau-1}(\omega))(\mathbf{P}_\tau(\omega))\right) = \sum_{\tau=1}^{T} \lambda_t U_{\omega_0}\left(f(\mathbf{P}_{\tau-1}(\omega))(\mathbf{P}_\tau(\omega))\right)$.

Hence, it must be that $\mu_{\omega_0}(\omega_t)\mu_{\omega_0}(s|\omega_t) = Q(\{\omega|\mathbf{P}_{t+1}(\omega) = (\omega_t, s)\})$ and thus that $\pi_{\omega_0}(\omega) = Q(\omega)$. It follows that $\mu_{\omega_0}(\omega_t) = \pi_{\omega_0}(\{\omega|\mathbf{P}_t(\omega) = \omega_t)\}$.

A similar proof can be done for the utility function at each $\omega_t \in \mathbf{\Omega}_{T-1}$. Therefore, for any partial history, $\mu_{\omega_t}(\omega_\tau) = \pi_{\omega_t}(\{\omega|\mathbf{P}_\tau(\omega) = \omega_\tau)\}$.

Lemma 10. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$, and $\hat{c} \in \widehat{C}(\omega_\tau) \setminus C(\omega_t)$, $u^*_{\omega_t}(\omega_\tau)(\hat{c}) = u^*_{\omega_t}(x_{\omega_t})$.*

Proof of Lemma 10: By Axiom 6,

$$u^*_{\omega_t}(\omega_\tau)(\hat{c}) = u^*_{\omega_t}(\omega_\tau)(\breve{c}) \tag{53}$$

for all $\hat{c}, \breve{c} \in \hat{C}(\omega_\tau) \setminus C(\omega_t)$. Also by Axiom 6,

$$u^*_{\omega_t}(\omega_\tau)(x_{\omega_t}) = u^*_{\omega_t}(\omega_\tau)(x_{\omega_\tau}). \tag{54}$$

By Axiom 8(ii),

$$u^*_{\omega_t}(\omega_t)(x_{\omega_t}) = \alpha u^*_{\omega_t}(\omega_t)(c^*) + (1-\alpha) u^*_{\omega_t}(\omega_t)(c_*) \tag{55}$$

$$\Rightarrow u^*_{\omega_t}(\omega_\tau)(x_{\omega_t}) = \alpha u^*_{\omega_t}(\omega_\tau)(c^*) + (1-\alpha) u^*_{\omega_t}(\omega_\tau)(c_*) \tag{56}$$

By Lemma 4, $u^*_{\omega_t}(\omega_\tau)$ agrees with $u_{\omega_t}(\omega_\tau)$ on $C(\omega_t)$ for all $\omega_t \in \mathbf{\Omega}_{T-1}$ and $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. By Lemma 5, $u_{\omega_t}(\omega_\tau)(c) = u_{\omega_t}(c)$ for all $c \in C(\omega_t)$. Therefore, the right hand sides of (55) and (56) are equal, which implies that $u^*_{\omega_t}(\omega_\tau)(x_{\omega_t}) = u^*_{\omega_t}(\omega_t)(x_{\omega_t}) \equiv u^*_{\omega_t}(x_{\omega_t})$. Equation (54) now implies that $u^*_{\omega_t}(\omega_\tau)(x_{\omega_\tau}) = u^*_{\omega_t}(x_{\omega_t})$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ and (53) implies that $u^*_{\omega_t}(\omega_\tau)(\hat{c}) = u^*_{\omega_t}(x_{\omega_t})$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ and $\hat{c} \in \widehat{C}(\omega_\tau) \setminus C(\omega_t)$.

Lemma 11. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, $u^*_{\omega_t}(x_{\omega_t}) = u^*(x_{\omega_t}) \equiv u^*$.*

Proof of Lemma 11: By Lemmas 3, 4, and 10,

$$\begin{aligned} & x(C(\omega_t))_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)} f \sim_{\omega_t} (\alpha c^* + (1-\alpha)c_*)_{S_{t+1}(\omega_t)\setminus\tilde{S}_{t+1}(\omega_t)} f \\ \Leftrightarrow & u^*_{\omega_t}(x_{\omega_t}) = \alpha u_{\omega_t}(c^*) + (1-\alpha)u_{\omega_t}(c_*) \end{aligned} \tag{57}$$

and

$$\begin{aligned} & (x_{(\omega_t,s_{t+1})})_{S_{t+2}(\omega_t,s_{t+1})\setminus\tilde{S}_{t+2}(\omega_t,s_{t+1})} f \sim_{(\omega_t,s_{t+1})} (\alpha c^* + (1-\alpha)c_*)_{S_{t+2}(\omega_t,s_{t+2})\setminus\tilde{S}_{t+2}(\omega_t,s_{t+2})} f \\ \Leftrightarrow & u^*_{(\omega_t,s_{t+1})}(x_{(\omega_t,s_{t+1})}) = \alpha u_{(\omega_t,s_{t+1})}(c^*) + (1-\alpha)u_{(\omega_t,s_{t+1})}(c_*) \end{aligned} \tag{58}$$

By Lemma 6, $u_{(\omega_t,s_{t+1})}(c) = u_{\omega_t}(c)$ for all $c \in C(\omega_t)$. Thus, the right hand sides of (57) and (58) are equal. By Axiom 8(i), (57) implies (58). Thus, $u^*_{(\omega_t,s_{t+1})}(x_{(\omega_t,s_{t+1})}) = u^*_{\omega_t}(x_{\omega_t})$. One can proceed by induction to show that $u^*_{\omega_\tau}(x_{\omega_\tau}) = u^*_{\omega_t}(x_{\omega_t})$ for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$. Setting $t = 0$, it follows that

$$u^*_{\omega_\tau}(x_{\omega_\tau}) = u^*_{\omega_0}(x_{\omega_0}) \equiv u^*(x_{\omega_0}).$$

Since all other $c \in C(\omega_\tau)$ can be evaluated by u_{ω_τ}, x_{ω_τ} is the only 'consequence' that needs to be evaluated by $u^*_{\omega_\tau}$. Thus, one can define $u^* \equiv u^*(x_{\omega_0})$ and use u^* as a parameter in the representation.

Remark L11: The parameter u^* depends on which transformation of the Bernoulli-utility function we use. Setting $u^* = 0$ fixes the level of the Bernoulli utility function such that if u_{ω_0} and $\tilde{u}_{\omega_0}$ both represent $\succsim_{\omega_0}$, then $\tilde{u}_{\omega_0} = \alpha u_{\omega_0}$ for $\alpha > 0$.

Lemma 12. *For all $\omega_t \in \mathbf{\Omega}_{T-1}$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_t)$ and for all $s \in S_{t+1}(\omega_\tau)$, define $\pi_{\omega_t}(\omega_\tau, s) = \mu_{\omega_t}(\omega_\tau)\mu_{\omega_t}(s|\omega_\tau)$. Then*

$$\begin{aligned} V_{\omega_t}(f) = \sum_{\tau=t}^{T-1} \beta^\tau \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \sum_{s \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, s) \Bigg(& \sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c) u_{\omega_t}(c) \\ & + \Big(1 - \sum_{c \in C(\omega_t)} f(\omega_\tau)(s)(c)\Big) u^* \Bigg) \end{aligned} \tag{59}$$

Remark L12: Setting $u^* = 0$, (59) reduces to (15).

Proof of Lemma 12: This follows from Lemmas 1 through 11.

Lemma 13. *The probability measures π_{ω_t} satisfy that for all $\omega_{t+1} \in \Omega_{t+1}(\omega_t)$, for all $\omega_\tau \in \mathbf{\Omega}_{T-1}(\omega_{t+1})$, and for all $s, \tilde{s} \in S_{t+1}(\omega_t)$, we have that*

$$\frac{\pi_{\omega_t}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(s|\omega_\tau)\})}{\pi_{\omega_t}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)\})} = \frac{\pi_{\omega_{t+1}}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(s|\omega_\tau)\})}{\pi_{\omega_{t+1}}(\{\omega : \mathbf{P}_{\tau+1}(\omega) \in \mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)\})}.$$

Proof of Lemma 13: Let g and h be as in Axiom 9. Then

$$\begin{aligned} g \succsim_{\omega_t} h \Leftrightarrow & \sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, \tilde{s}) u_{\omega_t}(\eta c^* + (1-\eta) c_*) \\ & \geq \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) u_{\omega_t}(c^*) + \left(\sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, \tilde{s}) - \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) \right) u_{\omega_t}(c_*) \\ & \Leftrightarrow \sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, \tilde{s}) [u_{\omega_t}(\eta c^* + (1-\eta) c_*) - u_{\omega_t}(c_*)] \\ & \geq \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) [u_{\omega_t}(c^*) - u_{\omega_t}(c_*)], \end{aligned} \tag{60}$$

and

$$\begin{aligned} g \succsim_{\omega_{t+1}} h \Leftrightarrow & \sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_{t+1}}(\omega_\tau, \tilde{s}) [u_{\omega_{t+1}}(\eta c^* + (1-\eta) c_*) - u_{\omega_{t+1}}(c_*)] \\ & \geq \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) [u_{\omega_{t+1}}(c^*) - u_{\omega_{t+1}}(c_*)], \end{aligned} \tag{61}$$

By Lemma 6, $u_{\omega_{t+1}} = u_{\omega_t}$. Thus, when (60) and (61) hold with equality, they and Axiom 9 imply that

$$\frac{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, s)} = \frac{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_{t+1}}(\omega_\tau, s)}. \tag{62}$$

A relationship like the one in (62) holds for all states $s \in S_{t+1}(\omega_t)$. Therefore, we have the result in (16).

Proof of sufficiency of Axioms: The result follows from Lemmas 1 through 13.

A.2 Necessity of Axioms

Necessity of Axiom 1 is obvious. Necessity of Axiom 2 follows from the mixture space theorem. Necessity of Axiom 4 follows from u being non-constant and π_{ω_t} having full support on $\Omega_{T-1}(\omega_t)$.

Axiom 3 is necessary, since the utilities for states where the LHS and RHS acts agree cancel out and one can divide through with the probabilities so that the utilities reduce to the same expressions for the two rankings in the axiom. A similar argument shows necessity of Axiom 5. Necessity of Axiom 6 follows from all $\hat{c} \notin C(\omega_t)$ being assigned the same utility value $u^* = 0$. Axiom 8(i) follows from $u^* = 0$ being invariant to the partial history under evaluation ω_τ. For $T \geq 2$, Axiom 8(ii) follows from $u^* = 0$ being invariant to the awareness level ω_t.

To show necessity of Axiom 7 for $T \geq 2$, note that

$$\begin{aligned}
& p_\tau p'_{\tau+1} l \succsim_{\omega_{\hat{t}}} q_\tau q'_{\tau+1} l && (63)\\
\Leftrightarrow & V_{\omega_{\hat{t}}}(p_\tau p'_{\tau+1} l) \geq V_{\omega_{\hat{t}}}(q_\tau q'_{\tau+1} l)\\
\Leftrightarrow & \beta^{\tau-\hat{t}} \sum_{c \in C(\omega_t)} p(c) u_{\omega_{\hat{t}}}(c) + \beta^{\tau-\hat{t}+1} \sum_{c \in C(\omega_t)} p'(c) u_{\omega_{\hat{t}}}(c)\\
& \geq \beta^{\tau-\hat{t}} \sum_{c \in C(\omega_t)} q(c) u_{\omega_{\hat{t}}}(c) + \beta^{\tau-\hat{t}+1} \sum_{c \in C(\omega_t)} q'(c) u_{\omega_{\hat{t}}}(c)\\
\Leftrightarrow & \sum_{c \in C(\omega_t)} p(c) u_{\omega_{\hat{t}}}(c) + \beta \sum_{c \in C(\omega_t)} p'(c) u_{\omega_{\hat{t}}}(c) \geq \sum_{c \in C(\omega_t)} q(c) u_{\omega_{\hat{t}}}(c) + \beta \sum_{c \in C(\omega_t)} q'(c) u_{\omega_{\hat{t}}}(c) && (64)
\end{aligned}$$

For different $\omega_{\tilde{t}}, \omega_{\tilde{\tau}}$, it holds that

$$\begin{aligned}
& p_{\tilde{\tau}} p'_{\tilde{\tau}+1} l \succsim_{\omega_{\tilde{t}}} q_{\tilde{\tau}} q'_{\tilde{\tau}+1} l && (65)\\
\Leftrightarrow & V_{\omega_{\tilde{t}}}(p_{\tilde{\tau}} p'_{\tilde{\tau}+1} l) \geq V_{\omega_{\tilde{t}}}(q_{\tilde{\tau}} q'_{\tilde{\tau}+1} l)\\
\Leftrightarrow & \beta^{\tilde{\tau}-\tilde{t}} \sum_{c \in C(\omega_t)} p(c) u_{\omega_{\tilde{t}}}(c) + \beta^{\tilde{\tau}-\tilde{t}+1} \sum_{c \in C(\omega_t)} p'(c) u_{\omega_{\tilde{t}}}(c)\\
& \geq \beta^{\tilde{\tau}-\tilde{t}} \sum_{c \in C(\omega_t)} q(c) u_{\omega_{\tilde{t}}}(c) + \beta^{\tilde{\tau}-\tilde{t}+1} \sum_{c \in C(\omega_t)} q'(c) u_{\omega_{\tilde{t}}}(c)\\
\Leftrightarrow & \sum_{c \in C(\omega_t)} p(c) u_{\omega_{\tilde{t}}}(c) + \beta \sum_{c \in C(\omega_t)} p'(c) u_{\omega_{\tilde{t}}}(c) \geq \sum_{c \in C(\omega_t)} q(c) u_{\omega_{\tilde{t}}}(c) + \beta \sum_{c \in C(\omega_t)} q'(c) u_{\omega_{\tilde{t}}}(c) && (66)
\end{aligned}$$

Since $u_{\omega_{\hat{t}}}(c) = u_{\omega_{\tilde{t}}}(c)$ for all $c \in C(\omega_t)$ and for all $\omega_{\hat{t}}, \omega_{\tilde{t}} \in \mathbf{\Omega}_{T-1}(\omega_t)$, the expressions in (64) and (66) are equivalent, and the equivalence of (63) and (65) follows.

To show necessity of Axiom 9 for $T \geq 2$, note that

$$g \succsim_{\omega_t} h \Leftrightarrow \eta u_{\omega_t}(c^*) + (1-\eta) u_{\omega_t}(c_*) \geq \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) u_{\omega_t}(c^*) + (1 - \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) u_{\omega_t}(c_*),$$

which holds if and only if $\eta \geq \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))$. Also,

$$\begin{aligned}
g \succsim_{\omega_{t+1}} h \Leftrightarrow & \eta u_{\omega_{t+1}}(c^*) + (1-\eta) u_{\omega_{t+1}}(c_*)\\
& \geq \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) u_{\omega_{t+1}}(c^*) + (1 - \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))) u_{\omega_{t+1}}(c_*),
\end{aligned}$$

which holds if and only if $\eta \geq \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))$.

By (16),

$$\frac{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{1 - \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))} = \frac{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{1 - \pi_{(\omega_{t+1})}(\mathscr{E}_{\tau+1}(s|\omega_\tau))},$$

which is equivalent to $\pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau)) = \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))$. Hence, $\eta \geq \pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))$ if and only if $\eta \geq \pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))$, which establishes that Axiom 9 holds.

B Proof of Theorem 2

Since $u_{\omega_{t+1}}$ is an extension of u_{ω_t}, $u_{\omega_{t+1}}(c) = u_{\omega_t}(c)$ for all $c \in \Delta(C(\omega_t))$. Both V_{ω_t} and $V_{(\omega_t,s)}(f|C(\omega_t))$ assign the value zero to all $\hat{c} \notin \Delta(C(\omega_t))$ (note that x is one such $\hat{c}$). Thus, future lotteries that return different $\hat{c} \notin \Delta(C(\omega_t))$ with the same probabilities are equivalent from the ω_t-point of view for both the ω_t-utility function, V_{ω_t}, and the ω_t-forecast of the ω_{t+1}-utility function, $V_{(\omega_t,s)}(f|C(\omega_t))$. Therefore,

$$V_{(\omega_t,s)}(f|C(\omega_t)) = \sum_{\tau=t+1}^{T-1} \beta^{\tau-t-1} \sum_{\omega_\tau \in \Omega_\tau(\omega_t,s)} \sum_{\tilde{s} \in S_{\tau+1}(\omega_\tau,s)} \pi_{\omega_{t+1}}(\omega_\tau, \tilde{s}) U_{\omega_t}(f(\omega_\tau)(\tilde{s})).$$

Given the updating rule (16),

$$\frac{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\pi_{\omega_t}(\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau))} = \frac{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(s|\omega_\tau))}{\pi_{\omega_{t+1}}(\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau))}. \tag{67}$$

for all future events $\mathscr{E}_{\tau+1}(s|\omega_\tau)$ and $\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)$ defined in (14). By definition, the events $\mathscr{E}_{\tau+1}(s|\omega_\tau)$ and $\mathscr{E}_{\tau+1}(\tilde{s}|\omega_\tau)$ are measurable with respect to the current set of extended consequences $\widehat{C}(\omega_t)$. Thus, (67) implies that prior and posterior conditional subjective probabilities agree. Therefore,

$$\begin{aligned} V_{\omega_t}(f) &= \sum_{\tau=t}^{T-1} \beta^{\tau-t} \sum_{\omega_\tau \in \Omega_\tau(\omega_t)} \sum_{s \in S_{\tau+1}(\omega_\tau)} \pi_{\omega_t}(\omega_\tau, s) U_{\omega_t}(f(\omega_\tau)(s)) \\ &= \sum_{s \in S_{t+1}(\omega_t)} \pi_{\omega_t}(\omega_t, s) \left[U_{\omega_t}(f(\omega_t)(s)) + \beta V_{(\omega_t,s)}(f|C(\omega_t)) \right]. \end{aligned} \tag{68}$$

References

[1] Ahn, David S. and Haluk Ergin (2010) "Framing Contingencies," *Econometrica* 78, 655-695.

[2] Alon, Shiri (2015) "Worst-Case Expected Utility," *Journal of Mathematical Economics* 60, 43-48.

[3] Anscombe, Francis J. and Robert J. Aumann (1963) "A Definition of Subjective Probability," *Annals of Mathematical Statistics* 34, 199–205.

[4] Dietrich, Franz (2018) "Savage's Theorem under changing awareness," *Journal of Economic Theory* 176, 1-54.

[5] Dominiac, Adam and Gerelt Tserenjigmid (2020) "Ambiguity under Growing Awareness," Working Paper, Virginia Tech.

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

[7] Easley, David and Aldo Rustichini (1999) "Choice Without Beliefs," *Econometrica* 67, 1157-1184.

[8] Epstein, Larry G. and Martin Schneider (2003) "Recursive Multiple-Priors," *Journal of Economic Theory* 113, 1-31.

[9] Grant, Simon, Idione Meneghel, and Rabee Tourky (2017) "Learning under unawareness,"Working Paper, Australian National University.

[10] Grant, Simon and John Quiggin (2013a) "Inductive Reasoning about Unawareness," *Economic Theory* 54, 717–755.

[11] Grant, Simon and John Quiggin (2013b) "Bounded awareness, heuristics and the Precautionary Principle," *Journal of Economic Behavior and Organization* 93, 17–31.

[12] Grant, Simon and John Quiggin (2015) "A Preference Model for Choice Subject to Surprise," *Theory and Decision* 79, 167-180.

[13] Halpern, Joseph Y., Nan Rong, and Ashutosh Saxena (2010) "MDPs with Unawareness," *Proceedings of the Twenty-Sixth Conference on Uncertainty in AI* 2010, 228-235.

[14] Karni, Edi, Quitzé Valenzuela-Stookey, and Marie-Louise Vierø (forthcoming) "Reverse Bayesianism: A generalization," *The B. E. Journal of Theoretical Economics.*

[15] Karni, Edi and Marie-Louise Vierø (2013) ""Reverse Bayesianism": A Choice-Based Theory of Growing Awareness," *American Economic Review* 103, 2790–2810.

[16] Karni, Edi and Marie-Louise Vierø (2015) "Probabilistic Sophistication and Reverse Bayesianism," *Journal of Risk and Uncertainty* 50, 189-208.

[17] Karni, Edi and Marie-Louise Vierø (2017) "Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance," *Journal of Economic Theory* 168, 301-328.

[18] Kochov, Asen (2018) "A Behavioral Definition of Unforeseen Contingencies," *Journal of Economic Theory* 175, 265-290.

[19] Koopmans, Tjalling C. (1972) "Representation of preference orderings over time," in: C.B. McGuire, R. Radner (Eds.), *Decision and Organization: A Volume in Honor of Jacob Marschak*, North-Holland, Amsterdam, 79–100.

[20] Lehrer, Ehud and Roee Teper (2014) "Extension Rules or What Would the Sage Do?" *American Economic Journal: Microeconomics* 6, 5-22.

[21] Li, Jing (2008) "A Note on Unawareness and Zero Probability," Working Paper, University of Pennsylvania.

[22] Lucas, Robert E. (1978) "Asset Prices in an Exchange Economy," *Econometrica* 46, 1429-1445.

[23] Ozbay, Erkut Y. (2007) "Unawareness and Strategic Announcements in Games with Uncertainty," in Dov Samet, ed. *Proceedings of Theoretical Aspects of Rationality and Knowledge XI*, Presses Universitaires De Louvain, 231-238.

[24] Piermont, Evan (2017) "Introspective Unawareness and Observable Choice," *Games and Economic Behavior* 106, 134-152.

[25] Sargent, Thomas J. (1987) *Dynamic Macroeconomic Theory.* Harvard University Press, Cambridge, Massachusets.

[26] Savage, Leonard J. (1954) *The Foundations of Statistics.* Wiley, New York.

[27] Schipper, Burkhard C. (2013). "Awareness-Dependent Subjective Expected Utility," *International Journal of Game Theory* 42, 725-753.

[28] Walker, Oliver and Simon Dietz (2011) "A Representation Result for Choice under Conscious Unawareness," Working Paper, London School of Economics and Political Science.

[29] Walley, Peter (1996) "Inferences from Multinomial Data: Learning about a Bag of Marbles," *Journal of the Royal Statistical Society Series B* 58, 3-57.

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