A Definition of Unforeseen Contingencies*

Asen Kochov†

February 15, 2011

Abstract

The paper models an individual who may not foresee all relevant aspects of an uncertain environment. The model is axiomatic and provides a novel choice-theoretic characterization of the collection of foreseen events. It is proved that all recursive, consequentialist models imply perfect foresight and thus cannot accommodate unforeseen contingencies. In particular, the model is observationally distinct from recursive models of ambiguity. The process of learning implied by dynamic behavior generalizes the Bayesian model and permits the collection of foreseen events to expand over time.

*This work has benefited from the helpful comments of Paulo Barelli and Bart Lipman. I am especially grateful to my advisor Larry Epstein for his support and for many insightful discussions.

†Department of Economics, University of Rochester, Rochester, NY 14627. E-mail: akochov@mail.rochester.edu.

1 Introduction

1.1 Objectives

At least since Williamson [17], economists have argued that unforeseen contingencies play a central role dynamic programming and in explaining many of the contractual arrangements and economic institutions that we observe. Despite their apparent importance and the general intuition that no decision maker can envisage all relevant aspects of even the most restricted economic problems, the empirical implications of unforeseen contingencies remain largely unclear.

The objective of this paper is to propose a behavioral definition of unforeseen contingencies. The explicit choice-theoretic characterization ensures verifiability given suitable choice data and thus the potential empirical relevance of unforeseen events. Unlike previous work on the subject, identification is made possible by the adoption of a multi-period domain of choice. Notably, the importance of the domain was anticipated by Kreps [13, p. 278] who argued that intertemporal choice may provide a foundation for disentangling the empirical implications of unforeseen contingencies *vis-a-vis* existing models of choice.

A key result in the paper shows that practically all recursive models of dynamic choice imply perfect foresight and hence cannot accommodate unforeseen contingencies. The intuition behind the result is simple. In any recursive model, an optimal course of action can be computed by applying the method of backward induction. Evidently, backward induction presupposes that the individual can reason through every node in the decision tree. The result provides one way to formalize this intuition and attests to the adequacy of the definition.

The prevalence of recursive methods raises the question if unforeseen contingencies can be tractably and parsimoniously modeled. The paper proceeds axiomatically in order to develop a complete model of dynamic choice which allows for previously unanticipated events to take place. To ensure tractability, a minimal departure from expected utility is sought. The scope of the departure is guided by two main desiderata. First, a model of unforeseen contingencies must give a precise behavioral expression of what it means for an individual to be aware that there are contingencies he does not foresee. Such awareness seems crucial to applications. As Kreps [13, p. 259] argues, "if one imagines that neoclassical contractual forms - relational contracting for example - are adopted consciously by the transacting parties [...], then it seems patent that those individuals are acting in anticipation of unforeseen contingencies."

A second desideratum is that the model achieves a clear separation between the individual's 'perception of uncertainty' and his 'tastes'. Arguably, one of the most interesting aspects of unforeseen contingencies is how individuals revise their perception of the environment in response to objective information about the state of the world. In Savage, perception is defined by the individual's subjective prior and dynamic consistency delivers a compelling normative argument for Bayesian updating. Ideally, a model of unforeseen contingencies would derive an updating rule for both foresight and beliefs from an axiom that preserves the essential normative appeal of dynamic consistency.

The aforementioned result shows that dynamic consistency must be necessarily weakened in order to accommodate unforeseen contingencies. A natural and weaker assumption is to require that violations of dynamic consistency arise only when previously unanticipated events alter the individual's perception. As the paper shows, the axiom characterizes a generalization of Bayesian updating known in statistics as the method of retrospective conditioning. Diaconis and Zabell [4] provide a discussion and Zabell [18] emphasizes the importance of unanticipated events in problems of statistical inference. A main feature of the rule is that the collection of foreseen events expands over time.

1.2 Examples

A major challenge in identifying the empirical implications of limited foresight is outlined in the concluding remarks of Dekel, Lipman, and Rustichini [3, p.540]:

> "One stumbling block [...] is that we want to distinguish between unforeseen contingencies and 'standard' uncertainty aversion. By the latter, we mean those models, such as nonadditive probability, which are intended to represent an agent who knows the state space but not the appropriate probabilities and behaves 'conservatively' because of this lack of knowledge. Conceptually,

> at least, there is a difference between this and not knowing the state space and behaving conservatively as a result. However, this distinction is difficult to make precise. Intuitively, we want a problem in which it is difficult for the agent to translate the options into utilities and where the agent exhibits some aversion to this difficulty. But how can we distinguish this from translating the options easily into utility space but then being uncertainty averse on the utility space?"

The empirical implications of limited foresight and uncertainty aversion would be disentangled if one could find a class of actions whose ranking does not require an assessment of likelihoods but is sufficiently rich to reveal the individual's knowledge of the state space. Since beliefs would be irrelevant for the evaluation of such actions, 'conservative' behavior can be justified only if some contingencies are unforeseen. In effect, the 'stumbling block' alluded to by Dekel, Lipman, and Rustichini [3] reduces to the problem of finding such a class. The paper proposes one possible solution. As the next example shows, the adoption of a multi-period domain of choice plays a key role.

1.2.1 Static Behavior

Consider an individual, Jones, who is contemplating an investment in a foreign country. He has two actions at his disposal, L and R, whose consequences unfold over a three-period horizon. The relevant environment is described by two contingencies, 'storm' (S) and 'political unrest' (U). Uncertainty resolves successively and can be represented diagrammatically as the event tree in Fig. 1. Every branch of the tree corresponds to a complete resolution of uncertainty or a possible state of the world. For example, the uppermost branch corresponds to the state (S,U). In the figure, the event tree is depicted twice, once for each action, and numbers inside the boxes denote each action's respective state-contingent payoffs. For example, the choice of action L (left) results in a loss of *$2* in period $t = 2$ if the state (S,U) obtains. Finally, notice that no relevant events take place in the final period. This normalization is not essential to understanding the example and its significance would be clarified in later chapters.

The depicted representation of the actions L and R as *acts*, that is, as mappings from states of the world into payoffs, is interpreted as a descrip-

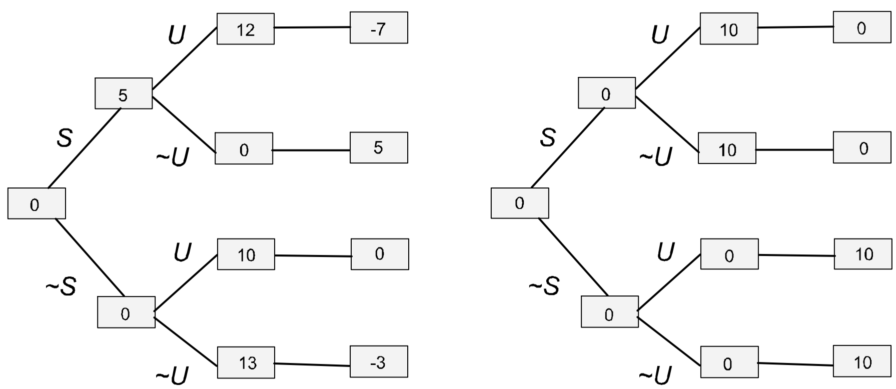

Figure 1: State-contingent payoffs of the actions L (left) and R (right).

tion of objective reality. Jones's perception is not a primitive of the model. Rather, in a manner consistent with revealed preference analysis, Jones's conceptualization of what might happen is derived endogenously given his observed ranking of the available actions. To keep details at a minimum and bring out the importance of the domain, suppose Jones is risk-neutral and does not discount the future.[1].

Suppose first that Jones 'foresees' all possible contingencies where foresight is to be understood informally. For example, one can imagine that Jones himself can draw the picture and sketch both physical actions as the exact state-contingent acts depicted in the figure. In the words of Dekel, Lipman, and Rustichini [3], Jones has perfect foresight if he can 'translate' all actions into 'utilities'. If Jones can do this, he can easily deduce that both actions are *effectively certain:* in every state of the world, that is, along every branch of the tree, the payoffs of either action add up to a total of *$10*. Jones would then be indifferent between L and R and the action which pays *$10* at the time of choice:

$$L \sim_0 10 \text{ and } R \sim_0 10 \tag{1.1}$$

The rankings in (1.1) are important in that they provide a refutable prediction of perfect foresight that is independent of whatever beliefs Jones might

[1] The assumption that the individual does not discount the future is relaxed in the formal model. For ease of exposition, risk neutrality is maintained throughout. It should be viewed as a normalization that can be justified in the familiar way by adopting an Anscombe and Aumann [2] formulation.

happen to entertain. In particular, these rankings obtain necessarily even if Jones was unsure about the likelihood of some events and was averse to the implied ambiguity. The point is that, so long as Jones foresees the events, his assessments of likelihood are irrelevant to the evaluation of effectively certain actions.

To turn these observations into an operational choice-theoretic definition of foresight, consider the rankings:

$$L \prec_0 10 \text{ and } R \sim_0 10 \tag{1.2}$$

How can one interpret such behavior? The effectively certain action R entails a ten-dollar bet on the event S paying in period $t = 2$; and exactly the same bet on ¬S paying in period $t = 3$. The indifference $R \sim_0 10$ reveals that Jones understands the 'structure' of the action R, namely, that the action comprises two bets which jointly offset one another. This could mean that Jones himself can draw a tree that looks very much like the state-contingent act in Fig. 1. Notice, however, that for the above indifference to obtain, Jones's tree need only depict the possibility of a storm. Mathematically, all Jones needs to do is draw the smallest tree with respect to which the act induced by R is measurable. To summarize, if such indifference obtains for effectively certain actions involving bets on the S and ¬S, the paper concludes that the events are *foreseen*.

In contrast, the second ranking $L \prec_0 10$ reveals that Jones cannot translate the action L into the finely-tuned act depicted in the figure. Such knowledge would be sufficient to deduce the effective certainty of the action. Such knowledge seems to be also necessary. For example, thinking of his possible payoffs in the event of no storm, Jones might have some understanding that his payoffs are uncertain. He may further know that he can win at least *$10* in period $t = 2$, and lose at most *$3* in the next. However, without foreseeing the exact tree, it is difficult to deduce that the implicit uncertainty cancels out. The ranking $L \prec_0 10$ is thus interpreted as being indicative of the fact that political unrest is unforeseen.

The ability to test Jones's knowledge of the state space independently of his beliefs is specific to the multi-period domain of choice and the existence of nontrivial effectively certain actions. If all uncertainty resolves in a single period, the only effectively certain actions are constant acts whose ranking reveals no information about Jones's perception of events. Conversely, any nontrivial act is uncertain and its evaluation requires an assessment of beliefs.

This is the 'stumbling block' encountered in atemporal settings and alluded to in Dekel, Lipman, and Rustichini [3]. Put differently, the atemporal domain is too 'small' to separate ambiguity aversion from limited foresight. The conclusion is confirmed by the results in Epstein, Marinacci, and Seo [7], Ghirardato [9], Gilboa and Schmeidler [11] and Mukerji [15].

1.2.2 Dynamic Behavior

To illustrate the implications of unforeseen contingencies for dynamic behavior, suppose Jones has limited foresight in period $t = 0$ but understands the environment perfectly in period $t = 1$. Denote his conditional preferences in the latter period by $\succeq_{1,S}$ and $\succeq_{1,\neg S}$, as the event S, or respectively ¬S, obtains. The choices below contrast Jones's posterior and prior valuations of the actions L and R:

$$L \sim_{1,S} 10, \qquad L \sim_{1,\neg S} 10 \quad \text{and} \quad L \prec_0 10 \tag{1.3}$$

$$R \sim_{1,S} 10, \qquad R \sim_{1,\neg S} 10 \quad \text{and} \quad R \sim_0 10 \tag{1.4}$$

The rankings in (1.3) and (1.4) reveal two important characteristics of dynamic choice under limited foresight.

The premium required for distant, poorly foreseen bets disappears as time unfolds and Jones's perception of the environment improves. The corresponding rankings in (1.3) imply a violation of dynamic consistency that is precluded by the standard approach to dynamic choice. The violation arises as Jones learns aspects of the environment he did not anticipate and could not have taken into account ex ante.

In contrast, the indifference $R \sim_0 10$ reveals that Jones foresees the immediate possibility of a storm. Then, he evaluates the action consistently over time as indicated by the conditional preferences in (1.4). The rankings illustrate the paper's approach to modeling coherent dynamic behavior when some contingencies are unforeseen: The individual is forward-looking and revises his plans *only when* unanticipated circumstances contradict his perception.

2 Static Model

2.1 Domain

The objective environment is described by a state space Ω and a finitely generated filtration $\mathcal{F} := \{\mathcal{F}_t\}$ where time varies over a finite horizon $\mathcal{T} = \{0, 1, ..., T\}$. An action taken by the individual induces a real-valued, $\{\mathcal{F}_t\}$-adapted process of outcomes. Call any such process an *act* and denote generic acts by h, h'. An act h is often written as a sequence $(h_1, h_2, ...h_T)$ where $h_t : \Omega \to \mathbb{R}$ is $\mathcal{F}_t$-measurable for every t. The set of all acts $\mathcal{H}$ is a mixture space under the obvious operation.

The distinction between *actions* and *acts* is an important part of the model. Physical actions, such as the purchase of a dividend-paying stock, comprise the individual's domain of choice. Acts are a mathematical construct used to model the relevant uncertainty: each act maps states of the world into outcomes. The paper assumes that the modeler observes the ranking of actions and constructs the corresponding acts. Under the assumption that each act is induced by a unique action, the observable choice over actions induces a unique preference over acts. The latter is adopted as a primitive of the model.

For any act h, let $\mathcal{F}(h)$ denote the smallest filtration with respect to which the act h is adapted. Conversely, for any subfiltration $\mathcal{G}$ of $\mathcal{F}$, $\mathcal{H}_{\mathcal{G}}$ is the subset of all $\mathcal{G}$-adapted acts. For every t, $\Pi_{\mathcal{G}_t}$ is the partition generating the algebra $\mathcal{G}_t$ and, for every ω, $\mathcal{G}_t(\omega)$ denotes the atom in $\Pi_{\mathcal{G}_t}$ containing ω. Since $\mathcal{F}_T$ is finitely generated, the latter is well-defined. It is assumed that $\mathcal{F}_T = \mathcal{F}_{T-1}$. Thus, the individual lives for another period after all relevant uncertainty is resolved. As explained in the next section, the assumption guarantees that the subset of effectively certain acts is rich. In particular, it generates the filtration $\mathcal{F}$.[2]

For any outcome $x \in \mathbb{R}$, let $\mathbf{x}$ denote the constant act paying x in every period and every state of the world. Given any act h and state ω, let $h(\omega)$ denote the act $(h_1(\omega), h_2(\omega)..., h_T(\omega))$. Note that outcomes of the act $h(\omega)$ may depend on the time period but not on the state of the world.

[2] The assumption is redundant in an infinite-horizon model. See the next section for further discussion.

2.2 Definition of Foreseen Events

The section introduces the behavioral definition of foreseen events. The explicit, choice-theoretic characterization ensures verifiability given suitable choice data and thus the potential empirical relevance of unforeseen contingencies. The main preliminary step identifies a rich class of acts whose ranking does not require an assessment of likelihoods but is sufficiently rich to reveal the individual's knowledge of the environment.

Building on the introductory example, define an act h to be **effectively certain** if

$$h(\omega) \sim h(\omega') \text{ for all } \omega, \omega' \in \Omega.$$

Theorem 6 shows that any effectively certain act h is necessarily indifferent to the act $h(\omega)$ whenever preference can be represented by a recursive utility function. The motivating example suggests that such indifference is intuitive *only if* the individual has perfect knowledge of the relevant environment. Consequently, and in contrast to the standard model, the paper does not impose indifference for all effectively certain acts *a priori*. Instead, it takes the subset of acts for which indifference obtains as *indicative* of the collection of foreseen events.

An effectively certain act h is **subjectively certain** if for all effectively certain, $\mathcal{F}(h)$-adapted acts h':

$$h' \sim h'(\omega) \text{ for all } \omega \in \Omega.$$

The definition of subjective certainty incorporates two complementary requirements. First, any subjectively certain act h must be indifferent to $h(\omega)$ for every $\omega \in \Omega$. This is true for the act R in the introductory example (1.2). There, the indifference $R \sim 10$ reveals that the Jones *foresees* the contingencies S and $\neg$S. Second, the subjective certainty of h requires that the same indifference obtain for all other $\mathcal{F}(h)$-adapted, effectively certain acts. For example, suppose one were to replace the *\$10* outcomes of R by *\$15.* The new act, say R', is effectively certain, and, by construction, entails bets on the same events S and $\neg$S. The original act R is then subjectively certain only if R' is indifferent to the sure payment of *\$15.* In effect, subjective certainty is a property of the *events* S and $\neg$S and is robust to changes in the act's outcomes. The property motivates the behavioral definition of foreseen events.

Definition 1 *An event is **foreseen** if it belongs to the filtration, $\mathcal{G}$, induced by the subset of subjectively certain acts. An act is **foreseen** if it is $\mathcal{G}$-adapted.*

The scope and applicability of this definition merit some discussion. As a behavioral test of unforeseen contingencies, subjective certainty is powerful *only* in a multi-period domain of choice. If all uncertainty resolves in a single period, the ranking $h \not\sim h(\omega)$ may be alternatively interpreted as an instance of state-dependent preferences. This interpretation is *not* valid in a temporal setting because subjective certainty is *no longer* implied by state-independence. The latter is evident from the ranking $L \prec 10$ in (1.2) and the fact that preferences in the introductory example are risk-neutral and *a fortiori* state-independent.

The temporal domain is important in that it permits the construction of nontrivial effectively certain acts. For example, consider the construction of the act R: for a given event A, there is a *$10* bet on A that pays off in period k, and an analogous bet on A^c that pays off in period $k' \neq k$. Thus, the event A belongs to at least two algebras $\mathcal{F}_k \neq \mathcal{F}_{k'}$ in the objective event tree $\{\mathcal{F}_t\}$. In a finite-horizon model, this is possible if and only if $A \in \mathcal{F}_{T-1}$. Consequently, the proposed distinction between 'foreseen' and 'unforeseen' is empirically relevant only for events in $\mathcal{F}_{T-1}$. In particular, the distinction becomes moot in the extreme case when $\mathcal{F}_0 = \mathcal{F}_{T-1}$. The latter case is isomorphic to an atemporal model since all uncertainty resolves in a single period. For expositional ease, the paper assumes that $\mathcal{F}_T = \mathcal{F}_{T-1}$, or equivalently, that the distinction is relevant for all events in the objective environment.

The subset of effectively certain actions is the *smallest* subset whose ranking is *sufficient* to identify the collection of foreseen events. The empirical content of limited foresight, however, is not exhausted by the specification of foreseen events. How does a self-aware individual perceive and evaluate actions whose outcomes depend on unforeseen states of the world? How do we model learning in response to objective information? The second question is especially pertinent since limited foresight *necessitates* a weakening of dynamic consistency and thus a generalization of Bayesian updating. An answer to these questions requires that the definition is incorporated in a complete model of dynamic choice. The rest of the paper proceeds axiomatically to develop and characterize such a model.

2.3 Axioms

The individual's choices among physical actions induce a preference ordering $\succeq$ on the set of objective acts $\mathcal{H}$. This section adopts a set of axioms on the induced preference $\succeq$. Some of the axioms have a standard interpretation if the individual has perfect knowledge of the environment. In this case, his subjective perception and the primitive objective environment coincide. If perception is coarse, however, the axioms make implicit assumptions about how perception differs from, and approximates, the objective world. The first two axioms fall in this category and their content is reinterpreted accordingly.

Order *The preference $\succeq$ is complete and transitive.*

Mixture Continuity For all acts $h \succ h' \succ h''$, there exist $\alpha, \beta \in (0,1)$ such that

$$h \succ \alpha h + (1-\alpha)h'' \succ h' \succ \beta h + (1-\beta)h''$$

Convexity *For all $h, h' \in \mathcal{H}$, $h \sim h'$ implies $\alpha h + (1-\alpha)h' \succeq h$.*

To understand Convexity, it is useful to imagine a hypothetical, auxiliary step in which the individual is asked to compare the *subjective mixture* of h and h' to either action. That is, the individual can mix the outcomes of h and h' as *he perceives* them. The subjective mixture 'smooths' outcomes across states foreseen by the individual. Aware that his perception may be incomplete, the individual prefers the mixture. The latter hedges his exposure to contingencies that he fears might be only a coarse approximation to the world. Convexity requires that the objective mixture $\alpha h+(1-\alpha)h'$ be preferred to its subjective counterpart. The former smooths the uncertainty *within* as well as *across* any of the foreseen events.

Certainty Independence *For all acts $h, h' \in \mathcal{H}$ and effectively certain, foreseen acts g:*

$$h \succeq h' \text{ if and only if } \alpha h + (1-\alpha)g \succeq \alpha h' + (1-\alpha)g.$$

To understand Certainty Independence, imagine the mixture of the action L and the subjectively certain action R in examples (1.1) and (1.2). Since R is foreseen, it is constant within the events S and ¬S foreseen by the individual. As such, it cannot hedge the poorly understood uncertainty *within* these events. Since the action is effectively certain, it also provides no hedging *across* the collection of foreseen events. The conjunction of these arguments motivates Certainty Independence.

To illustrate the next axiom, consider the foreseen action R in the introductory example and the corresponding ranking $R \sim 10$. The action R entails a ten-dollar bet on the event S paying in period $t = 2$ and an analogous bet on the event ¬S paying in period $t = 3$. Suppose one were to delay the payment of these bets by one period. The new action entails bets on the same events S and ¬S but pays in periods $t = 3$ and $t = 4$, respectively. Since the events are foreseen, Stationarity requires that the ranking remains the same.

Stationarity *For all acts $h, h' \in \mathcal{H}$ and for all outcomes $x \in M$,*

$$(h_0, ..., h_{T-1}, x) \succeq (h'_0, ..., h'_{T-1}, x) \text{ if and only if } (x, h_0, ..., h_{T-1}) \succeq (x, h'_0, ..., h'_{T-1}),$$

whenever the acts on the left (right) are foreseen.

The set of nodes $\cup_t \Pi_{\mathcal{G}_t}$ in the filtration of foreseen events correspond to states of the world as *perceived* by the individual. The next axiom is a subjective analogue of the standard monotonicity or state-independence assumption applied to these subjective states.

Subjective Monotonicity *For all acts $h, h' \in \mathcal{H}$ and for all outcomes $x \in \mathbb{R}$,*

$$hAx \succeq h'Ax \text{ for all } A \in \cup_t \Pi_{\mathcal{G}_t} \text{ implies } h \succeq h'.$$

Consider an action h whose continuation act at some node $\mathcal{F}_t(\omega)$ is non-constant: $h_\tau(\omega') \neq h_\tau(\omega'')$ for some $\tau > t$ and $\omega', \omega'' \in \mathcal{F}_t(\omega)$. Say that h' *simplifies* h if h' has a constant continuation at $\mathcal{F}_t(\omega)$ and equals h elsewhere. By construction, h' depends on events that are strictly closer in time. The next axiom requires that h' is foreseen whenever h is. Thus, events closer in time are easier to foresee.

Simplification *If g' simplifies a foreseen act g, then g' is foreseen.*

Define nullity in the usual way: the event $A \in \mathcal{F}_T$ is $\succeq$-null if $h(\omega) = h'(\omega)$ for all $\omega \in A^c$ implies $h \sim h'$. The next axiom ensures that every foreseen event is nonnull.

Nonnullity *For every foreseen event A, there exist outcomes $x > y, z$* such that $xAz \succ yAz$.

2.4 Representation

2.4.1 The Subjective Filtration

The section introduces the class of filtrations used to model the individual's perception of the objective environment $(\Omega, \{\mathcal{F}_t\})$.

Definition 2 *A subfiltration $\{\mathcal{G}_t\}$ of $\{\mathcal{F}_t\}$ is **sequentially connected** if*

$$\Pi_{\mathcal{G}_t} \backslash \Pi_{\mathcal{F}_t} \subset \Pi_{\mathcal{G}_{t+1}} \text{ for all } t < T.$$

To understand the definition, consider two disjoint events A_1 and A_2, either of which may be realized in period t. The individual foresees their union $A_1 \cup A_2$ but does not foresee the finer contingencies A_1 and A_2. Thus, $A_1 \cup A_2 \in \Pi_{\mathcal{G}_t} \backslash \Pi_{\mathcal{F}_t}$. The definition of a *sequentially connected filtration* requires that whenever the individual's perception of period t is coarse, he does not foresee any of the *more distant* contingencies within $A_1 \cup A_2$. In particular, $A_1 \cup A_2 \in \Pi_{\mathcal{G}_{t+1}}$. The requirement captures the intuition that events more distant in time are more difficult to foresee.

It is not difficult to see that any sequentially connected filtration $\{\mathcal{G}_t\}$ is fully determined by the algebra $\mathcal{G}_T$:

$$\mathcal{G}_t = \mathcal{F}_t \cap \mathcal{G}_T \text{ for all } t \in \mathcal{T}. \tag{2.1}$$

Define an algebra to be *sequentially connected*, if it induces a sequentially connected filtration via (2.1). The rest of the paper uses $\mathcal{G}$ interchangeably to denote the filtration and the algebra which generates it.

Sequentially connected filtrations include a number of special cases which admit intuitive interpretations.

Example 1 *(Fixed Horizon) The individual foresees all events up to some period k:*

$$\mathcal{G}_t = \mathcal{F}_t \text{ for all } t \leq k \text{ and } \mathcal{G}_t = \mathcal{F}_k \text{ for all } t > k.$$

More generally, the individual may not foresee the contingencies describing an unlikely event A, but have a better understanding of the contingencies describing its complement. his depth of foresight is then a random variable and the corresponding sequentially connected algebra can be modeled as a *stopping time*.

Example 2 *(Random Horizon) The individual foresees all events up to a stopping time τ, where*

$$\tau : \Omega \to \mathcal{T} \text{ and } \{\omega : \tau(\omega) = k\} \in \mathcal{F}_k \text{ for all } k \in \mathcal{T}.$$

The filtration $\{\mathcal{G}_t\}$ induced by the stopping time τ

$$\mathcal{G}_t := \{A \in \mathcal{F}_t : A \cap \{\omega : \tau(\omega) = k\} \in \mathcal{F}_k \text{ for all } k \in \mathcal{T}\} \text{ for } t \in \mathcal{T}$$

is sequentially connected.[3]

Sequentially connected filtrations arise as the outcome of a satisficing procedure for simplifying decision trees proposed by Gabaix and Laibson [8] in a setting of objective uncertainty.

Example 3 *(Satisficing) The individual ignores branches of the decision tree whose probability is lower than some threshold $\alpha \in [0, 1]$.*[4]

The Gabaix and Laibson [8] procedure permits a parsimonious parametrization of sequentially connected filtrations via the threshold parameter α.

The class of sequentially connected filtrations *excludes* the following subfiltration:

$$\mathcal{G} = \{\mathcal{F}_0, \mathcal{F}_0, ..., \mathcal{F}_0, \mathcal{F}_T\}$$

[3] Appendix 6.6 shows that sequentially connected filtrations inherit the lattice structure of stopping times. Specifically, the supremum (infimum) of sequentially connected filtrations is sequentially connected.

[4] Appendix 6.6 provides a detailed translation of the Gabaix-Laibson definition into the setting of this paper.

In this example, the individual *foresees all* possible contingencies ($\mathcal{G}_T = \mathcal{F}_T$) but he 'delays' the resolution of uncertainty. That is, he believes erroneously that all information is revealed in the last period. Since this filtration does not capture a coarse perception of the environment, it is precluded by Definition 2.

2.4.2 Subjective Acts

In Savage's [16] theory of subjective probability, the individual behaves as if he contemplates all possible states of the world and anticipates the outcomes that any given physical action might induce. That is, it is as though the individual himself perceives physical actions as the acts describing the relevant, objective uncertainty.

Savage [16, p. 92] was troubled by this implication of the theory and introduced the notion of *small worlds.* Like the subalgebra of foreseen events $\mathcal{G}$, a small world is a partition of the state space of all relevant uncertainty which Savage called the *grand world*. He used the following example. An individual (Jones) faces the decision whether to buy a certain convertible or not. In the simplest analysis, Jones perceives no uncertainty and regards the convertible as a sure enjoyment. Accordingly, Savage modeled Jones' perception of the convertible as a *small world consequence*. In reality, many contingencies may affect Jones' decision. As Savage wrote, Jones 'would not buy the convertible if he thought it likely that he would be immediately faced with a financial emergency arising out of the sickness of himself or of some member of his family.' The small world consequence is therefore only an approximation to a nonconstant *grand world act* in the universal state space.

Savage's primary emphasis in the *Foundations of Statistics* was normative. Both the small and grand worlds are therefore taken as primitive constructs which Jones could in principle use to make better decisions. From a descriptive standpoint, Jones' perception has to be derived *endogenously* from primitives observable to the modeler. To emphasize the endogeneity of small world acts, the paper refers to them as *subjective*. The modeler observes the ranking of *physical actions,* such as the convertible, and maps each physical action into a unique *act* in the set $\mathcal{H}$. The latter is interpreted as an *objective description* of the relevant environment in the sense that it is constructed by the modeler. These concepts and the relations between them can be represented diagrammatically:

$$
\begin{array}{ccc}
\text{Physical actions} & \overset{\text{1-1}}{\longleftrightarrow} & \text{Objective acts} \\
\widehat{\Phi} \;\downarrow & \swarrow \;\Phi & \\
\text{Subjective acts} & &
\end{array}
$$

The mapping $\widehat{\Phi}$ maps physical actions into *subjective acts*. The latter depict the individual's *perception* of the environment. The mapping Φ makes the graph commute.

The diagram makes it clear that even though physical actions are not explicitly modeled the individual's perception of actions can be represented by means of the mapping Φ from objective into subjective acts. Savage's theory of subjective expected utility can be interpreted as a special case in which Φ is the identity mapping from the set of objective acts $\mathcal{H}$ into itself. The next definition introduces a general class of mappings used to model Jones' perception of the convertible when some contingencies might be unforeseen.

Definition 3 *A $\mathcal{G}$-**approximation mapping** is a continuous, concave function Φ from $\mathcal{H}$ into $\mathcal{H}_{\mathcal{G}}$ such that for all acts $h, h' \in \mathcal{H}, g \in \mathcal{H}_{\mathcal{G}}$, nonnegative real numbers α and events $A \in \mathcal{G}_t$ for $t \in \mathcal{T}$:*

(i) $\Phi(\alpha h + g) = \alpha\Phi(h) + g$,

(ii) $(\Phi h)_t \mid A = (\Phi h')_t \mid A$ *whenever* $h_t \mid A = h'_t \mid A$.

Subjective acts respect the filtration $\mathcal{G}$ of foreseen events. Formally, any $\mathcal{G}$-approximation mapping takes every objective acts into $\mathcal{G}$-*adapted* subjective act. The property says that an individual cannot imagine bets on contingencies he does not foresee. In addition, property (i) implies that the function Φ maps each $\mathcal{G}$-*adapted* objective act into itself. Thus, the subjective and objective acts coincide for any action that is foreseen. The next example provides a simple, albeit extreme illustration of a $\mathcal{G}$-approximation mapping.

Example 4 *(Approximation From Below) For every $h \in \mathcal{H}$, let Φh be the lower $\mathcal{G}$-adapted envelope of h:*

$$(\Phi h)_t \mid A := \min_{\omega \in A} h_t(\omega), \text{ for every } t \in \mathcal{T} \text{ and } A \in \Pi_{\mathcal{G}_t}.$$

Then, Φ is a $\mathcal{G}$-approximation mapping.

The next lemma provides a convenient way to parametrize the class of $\mathcal{G}$-approximation mappings. For every t and $A \in \mathcal{F}_t$, let $ba_1(A, \mathcal{F}_t)$ denote the set of all additive set functions p such that $p(A) = 1$.

Lemma 1 *Φ is a $\mathcal{G}$-approximation mapping if and only if for every t and every $A \in \Pi_{\mathcal{G}_t}$ there exist a nonempty, closed, convex subset $\mathcal{C}_A$ of $ba_1(A, \mathcal{F}_t)$ such that:*

$$(\Phi h)_t \mid A = \min_{p \in \mathcal{C}_A} \int_A h_t \; dp. \tag{2.2}$$

2.4.3 Representation Theorem

This section completes the description of the static model.

Definition 4 *A preference relation $\succeq$ on $\mathcal{H}$ has a* ***limited foresight representation*** *$(\mathcal{G}, \Phi, \mathcal{C})$ if it admits a utility function of the form:*

$$V(h) = \min_{p \in \mathcal{C}} \int_\Omega \textstyle\sum_t \beta^t (\Phi h)_t \; dp, \tag{2.3}$$

where $\beta > 0$, $\mathcal{G}$ is the filtration of foreseen events, $\mathcal{G}$ is sequentially connected, Φ is a $\mathcal{G}$-approximation mapping, and $\mathcal{C}$ is a closed, convex subset in the interior of the simplex $\Delta(\Omega, \mathcal{G}_T)$.

The filtration $\mathcal{G}$ and the mapping Φ define the individual's model of the environment - the contingencies he foresees and the subjective act he attaches to any physical action. The nonsingleton set of priors $\mathcal{C}$ reflects the individual's awareness that his model of the environment may be incomplete.

Theorem 2 *A preference $\succeq$ satisfies Order, Mixture Continuity, Convexity, Certainty Independence, Stationarity, Subjective Monotonicity, Simplification and Nonnullity if and only if it has a limited foresight representation $(\mathcal{G}, \Phi, \mathcal{C})$. Moreover, the representation is unique.*

2.4.4 Discussion

In applications, one begins by first specifying the components of the representation: a sequentially connected filtration $\mathcal{G}'$, a mapping Φ' and a set of priors $\mathcal{C}'$ defined on the algebra $\mathcal{G}'_T$. Preference is then induced via the utility function in (2.3). A requirement of the representation is that $\mathcal{G}'$ is the filtration of foreseen events for the induced preference. The requirement is *not* satisfied for *any* choice of utility components. To emphasize the point, this section gives an example of a triple $(\mathcal{G}', \Phi', \mathcal{C}')$ for which $\mathcal{G}'$ is strictly coarser than the filtration of foreseen events. Sufficient conditions on the mapping Φ' are then provided that guarantee the required property and hence that $(\mathcal{G}', \Phi', \mathcal{C}')$ constitutes a limited foresight representation.

To begin, consider the triple $(\mathcal{G}', \Phi', \mathcal{C}')$, where $\mathcal{G}'$ is the trivial filtration $\{\mathcal{F}_0, \mathcal{F}_0, ..., \mathcal{F}_0\}$, $\mathcal{C}'$ is the degenerate measure on $\mathcal{F}_0 = \{\Omega, \varnothing\}$, and where the mapping Φ' is defined as follows:

$$(\Phi' h)_t(\omega) := \int h_t \; dp, \text{ for every } \omega \in \Omega \text{ and } t \in \mathcal{T}, \tag{2.4}$$

for some measure p on $\mathcal{F}_T$. The trivial filtration suggests the interpretation of zero foresight: the individual knows that 'something may happen' but is unable to specify any finer contingencies. Accordingly, the mapping Φ' takes each objective act into a *constant* subjective act.

The trivial filtration is sequentially connected and, by Lemma 1, the function Φ' is a $\mathcal{G}'$-approximation mapping. To verify whether $\mathcal{G}'$ is the collection of foreseen events, consider the preference ordering induced by the utility function in (2.3):

$$V'(h) := \textstyle\sum_t \beta^t (\Phi' h)_t. \tag{2.5}$$

Substituting (2.4) into (2.5) and rearranging, one obtains:

$$V'(h) = \textstyle\sum_t \beta^t \left(\displaystyle\int_\Omega h_t \; dp \right) = \displaystyle\int_\Omega \left(\textstyle\sum_t \beta^t h_t \right) \; dp. \tag{2.6}$$

The expected-utility representation in (2.6) shows that all events are foreseen in the sense of Definition 1. Conclude that the triple $(\mathcal{G}', \Phi', \mathcal{C}')$ is *not* a limited foresight representation of the induced preference. In fact, the *unique* such representation $(\mathcal{G}, \Phi, \mathcal{C})$ is given by the expected-utility functional in

(2.6) where $\mathcal{G}$ is the objective filtration $\mathcal{F}$, Φ is the identity mapping from $\mathcal{H}$ into itself, and beliefs $\mathcal{C}$ consist of the single prior p.

The example shows why an explicit choice-theoretic definition of foreseen events is important. For arbitrary triples $(\mathcal{G}', \Phi', \mathcal{C}')$, the intended distinction between 'foreseen' and 'unforeseen' events as embodied in the functional component $\mathcal{G}'$ may have no behavioral content. Moreover, for any given preference, there will be several candidate filtrations $\mathcal{G}'$ for which a utility exists and no guidance how to select the right one. A limited foresight representation makes such a selection. It has the merit of being based on an intuitive and behavioral characterization of foreseen events.

The next lemma proves that for *generic* triples $(\mathcal{G}', \Phi', \mathcal{C}')$, $\mathcal{G}'$ is in fact the filtration of foreseen events. To state the result, fix a sequentially connected filtration $\mathcal{G}'$ and a closed, convex set of priors $\mathcal{C}'$ in the interior of $\Delta(\Omega, \mathcal{G}'_T)$. Recall that, by Lemma 1, every $\mathcal{G}'$-approximation mapping Φ' can be identified with a collection $\{\mathcal{C}'_A\}$ where $\mathcal{C}'_A$ is a closed, convex subset of $ba_1(A, \mathcal{F}_t)$ as A and t vary over all cells of the filtration $\mathcal{G}'$.

Lemma 3 *A triple $(\mathcal{G}', \{\mathcal{C}'_A\}, \mathcal{C}')$ is a limited foresight representation of the induced preference whenever every set $\mathcal{C}'_A$ has nonempty interior.*

The following specifications of the approximation mapping satisfy the sufficient conditions.

Example 5 *(Approximation From Below) For every $t \in \mathcal{T}$, and $h \in \mathcal{H}$, $\Phi' h$ is the lower $\mathcal{G}'$-adapted envelope of h. By Lemma 3, $(\mathcal{G}', \Phi', \mathcal{C}')$ is a limited foresight representation of the induced preference.*

A notable feature of Example 5 is that the model is fully specified by the filtration $\mathcal{G}'$ and the set of priors $\mathcal{C}'$. A drawback is the 'coarseness' of the approximation mapping. In terms of the representation derived in Lemma 1, each set $\mathcal{C}'_A$ in the construction of Φ' equals the entire simplex $\Delta(A)$. To provide a less extreme approximation, take $\mathcal{C}'_A$ to be an ϵ-contraction of the simplex around a focal measure p^* in $\Delta(A)$:

$$\mathcal{C}'_A := \{\epsilon p + (1 - \epsilon) p^* : p \in \Delta(A)\}.$$

The corresponding approximation mapping is then fully determined by the filtration of foreseen events and the single parameter $0 < \epsilon < 1$.

Example 6 *(ϵ-Contamination) For every $t \in \mathcal{T}$ and every $A \in \Pi_{\mathcal{G}'_t}$, $\mathcal{C}_{A'}$ is an ϵ-contraction of the simplex. By Lemma 3, $(\mathcal{G}', \Phi'_\epsilon, \mathcal{C}')$ is a limited foresight representation of the induced preference.*

3 Dynamic Model

3.1 Axioms

This section develops the dynamic model of limited foresight. The primitive is an $\mathcal{F}$-adapted process of conditional preferences $\{\succeq_{t,\omega}\}$ where $\succeq_{t,\omega}$ describes the ranking of actions in state ω and period t. It is assumed throughout the section that every conditional preference $\succeq_{t,\omega}$ admits a limited foresight representation $(\mathcal{G}^{t,\omega}, \Phi^{t,\omega}, \mathcal{C}^{t,\omega})$.

The first axiom says that conditional preferences at t and ω do not depend on the actions' payoffs in any alternative history.

Consequentialism *For each t and ω and all acts h, h',*

$$h_\tau(\omega') = h'_\tau(\omega') \text{ for all } \tau \geq t \text{ and } \omega' \in F_t(\omega) \text{ implies } h \sim_{t,\omega} h'.$$

Consequentialism restricts the scope of limited foresight modeled in this paper. The stated indifference requires that whenever the relevant continuation acts are *objectively* identical, they are *perceived* as identical. Intuitively, perception may be incomplete, but not delusional. The individual does not imagine differences if there are none.

The next axiom requires that *tastes* do not depend on the time and state of the world. As in the static model, limited foresight pertains to the individual's perception of uncertainty and has no implications for the evaluation of constant acts.

Dynamic State Independence *For each t and ω,*

$$(x_0, ..., x_{t-1}, y_t, ..., y_T) \succeq_{t,\omega} (x_0, ..., x_{t-1}, y'_t, ..., y'_T) \text{ if and only if}$$

$$(x_0, ..., x_{t-1}, y_t, ..., y_T) \succeq_0 (x_0, ..., x_{t-1}, y'_t, ..., y'_T).$$

To understand the next axiom, consider an environment in which one of three events A_1, A_2 and A_3 can be realized in period $t = 1$. An individual with perfect foresight evaluates all actions consistently:

$$h \succeq_{1,A_i} h' \text{ for all } i \text{ implies } h \succeq_0 h'. \tag{3.1}$$

If he were to learn $A_1 \cup A_2$ at some hypothetical intermediate stage τ, then intertemporal consistency would similarly imply:

$$h \succeq_{1,A_1} h' \text{ and } h \succeq_{1,A_2} h' \text{ implies } h \succeq_{\tau,A_1\cup A_2} h'. \tag{3.2}$$

The implications in (3.1) and (3.2) reflect an individual who plans consistently. he knows all possible contingencies and anticipates accurately the future choices he is going to make. This knowledge is incorporated in her behavior in the period $t = 0$. It is important to emphasize that the hypothetical stage τ and the preferences $\succeq_{\tau,A_1\cup A_2}$, $\succeq_{\tau,A_3}$ in (3.2) are *not* part of the formal model. However, they prove useful in interpreting the implications of limited foresight for dynamic behavior.

To continue the example, consider an individual who initially foresees only the events $A_1 \cup A_2$ and A_3. Thinking of the future, he contemplates his behavior conditional on the events he foresees. If $\succeq^a_{1,A_1\cup A_2}$ and $\succeq^a_{1,A_3}$ denote these *anticipated preferences*, then:

$$h \ \succeq^a_{1,A_1\cup A_2} \ h' \text{ and } h \ \succeq^a_{1,A_3} \ h' \text{ implies } h \succeq_0 h'. \tag{3.3}$$

As in (3.1) and (3.2), the implication in (3.3) describes an individual who is forward-looking and plans ahead. However, the anticipated preferences reflect prior foresight and may differ from the individual's *actual* future behavior. Furthermore, the anticipated preferences may differ from the hypothetical rankings $\succeq_{\tau,A_1\cup A_2}$ and $\succeq_{\tau,A_3}$ used in (3.2). It turns out that the latter difference is easier to analyze since the conditioning events, namely, $A_1 \cup A_2$ and A_3, are the same in both cases. Specifically, the preference $\succeq^a_{1,A_1\cup A_2}$ represents behavior if the individual *were to learn* the event $A_1 \cup A_2$ and *perceive* the world as he does at $t = 0$. In contrast, the preference $\succeq_{\tau,A_1\cup A_2}$ represents behavior if the individual *were to learn* the event $A_1 \cup A_2$ and *perceive* the world as he does at $t = 1$.

These perceptions necessarily coincide only when the actions h and h' are foreseen at $t = 0$. Then (3.2) and (3.3) imply:

$$h \succeq_{1,A_i} h' \text{ for all } i \qquad \Rightarrow \qquad h \succeq_{\tau,A_1\cup A_2} h' \text{ and } h \succeq_{\tau,A_3} h'$$

$$\Rightarrow \qquad h \succeq^a_{1,A_1\cup A_2} h' \text{ and } h \succeq^a_{1,A_3} h'$$

$$\Rightarrow \qquad h \succeq_0 h'.$$

The above implications motivate the next axiom. It formalizes the present approach to modeling sophisticated dynamic behavior when some contingencies are unforeseen. Namely, it captures an individual who is forward-looking and revises his plans *only when* unanticipated circumstances contradict his perception. The approach is illustrated in the introductory examples (1.3) and (1.4).

Weak Dynamic Consistency *For each t and ω and for all acts g, g' in $\mathcal{H}_{\mathcal{G}^{t,\omega}}$ such that $g_\tau = g'_\tau$ for all $\tau \leq t$,*

$$g \succeq_{t+1,\omega'} g' \text{ for all } \omega' \text{ implies } g \succeq_{t,\omega} g'. \tag{3.4}$$

The standard axiom, *Dynamic Consistency*, requires that (3.4) obtain for all acts h, h' in $\mathcal{H}$ such that $h_\tau = h'_\tau$ *for all* $\tau \leq t$, irrespective of whether the acts are foreseen or not.

3.2 Representation

The process of learning implied by Consequentialism, State Independence and Weak Dynamic Consistency is derived. Specifically, the section characterizes how perception of the environment $\{\mathcal{G}^{t,\omega}, \Phi^{t,\omega}\}$ and beliefs $\{\mathcal{C}^{t,\omega}\}$ evolve over time.

The first implication captures a notion of expanding foresight. That is, for every t and ω, the posterior filtration $\mathcal{G}^{t+1,\omega}$ refines the prior filtration $\mathcal{G}^{t,\omega}$. To state this formally, let $\mathcal{G}^{t,\omega} \cap \mathcal{F}_{t+1}(\omega)$ denote the restriction of the prior filtration $\mathcal{G}^{t,\omega}$ to the subtree emanating from the node $\mathcal{F}_{t+1}(\omega)$.[5] The latter event is realized and is thus known by the individual at period $t+1$ and state ω.

Definition 5 *A process of filtrations $\{\mathcal{G}^{t,\omega}\}$ is* ***refining*** *if $\mathcal{G}^{t+1,\omega}$ refines $\mathcal{G}^{t,\omega} \cap \mathcal{F}_{t+1}(\omega)$ for all t and ω.*

[5] For every subset A of Ω, an algebra $\mathcal{G}$ on Ω induces the algebra $\mathcal{G}\cap A := \{B\cap A : B \in \mathcal{G}\}$ on A. A filtration $\{\mathcal{G}_t\}$ induces the filtration $\{\mathcal{G}_t\} \cap A := \{\mathcal{G}_t \cap A\}$ on A.

The next step describes how the process of conditional beliefs $\{\mathcal{C}^{t,\omega}\}$ evolves over time. Some preliminary definitions are necessary. For a set of priors $\overline{\mathcal{C}}$ on the objective algebra $\mathcal{F}_T$, define the set of prior-by-prior Bayesian updates by

$$\overline{\mathcal{C}}_t(\omega) := \{p(\cdot \mid \mathcal{F}_t(\omega)) : p \in \overline{\mathcal{C}}\},$$

and define the set of conditional one-step-ahead measures by

$$\overline{\mathcal{C}}_t^{+1}(\omega) := \{\text{marg}_{\mathcal{F}_{t+1}} p : p \in \overline{\mathcal{C}}_t(\omega)\}.$$

The following definition generalizes the familiar decomposition of a measure in terms of its conditionals and marginals to the decomposition of a set of measures $\overline{\mathcal{C}}$. The requirement is studied in Epstein and Schneider [5], who discuss its role for modeling dynamically consistent behavior when the individual has more than a single prior. Formally, define a set $\overline{\mathcal{C}}$ to be **$\{\mathcal{F}_t\}$-rectangular** if for all t and ω,

$$\overline{\mathcal{C}}_t(\omega) \;=\; \{\int_\Omega p_{t+1}(\omega')\; dm : p_{t+1}(\omega') \in \overline{\mathcal{C}}_{t+1}(\omega') \text{ for all } \omega' \text{ and } m \in \overline{\mathcal{C}}_t^{+1}(\omega)\}.$$

The main feature of rectangularity is that the decomposition on the right combines a marginal from $\overline{\mathcal{C}}_t^{+1}(\omega)$ with *any* measurable selection of conditionals. This will typically involve conditionals that are 'foreign' to a given marginal. If the set $\overline{\mathcal{C}}$ is a singleton, there are no foreign conditionals and the definition of rectangularity reduces to the standard decomposition of a probability measure.

Bayesian updating presupposes that prior beliefs are defined for every event in the objective algebra. Suppose instead that foresight expands gradually over time. This would be the case if, for example, a previously unanticipated event takes place. How should knowledge of the event be incorporated into the corpus of earlier beliefs? How should those beliefs be revised in response?

The rule to be described next is known in statistics as the method of *retrospective conditioning*. Diaconis and Zabell [4] provide a discussion. To convey its main idea, it is sufficient to focus on the special case when $\mathcal{C}^0$ consists of a single prior. The method comprises a two-step procedure. First, as the collection of foreseen events $\mathcal{G}^0$ expands over time, the individual *extends* his prior $\mathcal{C}^0$ to the new and finer algebra of foreseen events. The extension is *consistent* in that it preserves the original probabilities on $\mathcal{G}^0$. The extended prior is then updated by Bayes rule in response to objective information about the state of the world. The name 'retrospective' refers to the first stage of the method in which updating is temporarily suspended until a new 'prior' is 'retrospectively' quantified on the richer probability space $\mathcal{G}^1$. Afterwards, standard conditioning takes place.

The successive extension of beliefs results in a probability measure $\overline{\mathcal{C}}$ defined on the entire objective algebra $\mathcal{F}_T$. The method of retrospective conditioning can be alternatively and more parsimoniously formulated in terms of this 'shadow' probability. Namely, conditional beliefs $\mathcal{C}^{t,\omega}$ at every t and ω are obtained by (i) updating the measure $\overline{\mathcal{C}}$ using Bayes rule, and (ii) restricting the implied posterior to the algebra of foreseen events $\mathcal{G}^{t,\omega}$. This formulation is used in the definition below. In its general form, $\overline{\mathcal{C}}$ is a possibly nonsingleton, $\{\mathcal{F}_t\}$-rectangular set of measures and updating proceeds prior-by-prior.

Definition 6 *A process* $\{\mathcal{C}^{t,\omega}, \mathcal{G}^{t,\omega}\}$ ***admits a consistent extension*** *if there exists an* $\{\mathcal{F}_t\}$*-rectangular, closed and convex subset* $\overline{\mathcal{C}}$ *in the interior of* $\Delta(\Omega, \mathcal{F}_T)$ *such that*

$$\mathcal{C}^{t,\omega} = \{\text{marg}_{\mathcal{G}^{t,\omega}} p : p \in \overline{\mathcal{C}}_t(\omega)\} \text{ for every } t \text{ and } \omega.$$

The next theorem shows that Consequentialism, Temporal State Independence and Weak Dynamic Consistency fully characterize the process of learning described by Definitions 5 and 6.

Theorem 4 *The family of preferences* $\{\succeq_{t,\omega}\}$ *satisfies Consequentialism, Temporal State Independence and Weak Dynamic Consistency if and only if* $\{\mathcal{G}^{t,\omega}\}$ *is refining and* $\{\mathcal{C}^{t,\omega}, \mathcal{G}^{t,\omega}\}$ *admits a consistent extension* $\overline{\mathcal{C}}$*.*

The consistent extension $\overline{\mathcal{C}}$ is unique whenever the individual foresees all one-step-ahead contingencies. That is, for every t and ω, $\mathcal{F}_{t+1}(\omega)$ belongs to the collection of foreseen events $\mathcal{G}^{t,\omega}$.[6]

Theorem 5 *If* $\mathcal{F}_{t+1}(\omega) \in \mathcal{G}^{t,\omega}$ *for all* t *and* ω*, the consistent extension* $\overline{\mathcal{C}}$ *provided by Theorem 4 is unique.*

[6] Recall that $\mathcal{G}^{t,\omega}$ denotes both the sequentially connected filtration at t, ω and the algebra which generates it via (2.1).

4 Foresight and Dynamic Consistency

The section returns to the original problem of providing a choice-theoretic model-independent definition of foreseen events. Accordingly, a general setting is once again adopted whereby no utility representation is assumed for the process of conditional preferences $\{\succeq_{t,\omega}\}$. The analysis is motivated by Kreps' [13, p.278] intuition that dynamic behavior may reveal the collection of foreseen events and provide a foundation for separating limited foresight from existing models of uncertainty. The main result establishes a close connection between the intertemporal consistency of behavior and the static ranking of effectively certain actions. As a corollary, the section shows how intertemporal consistency may provide an alternative and equivalent characterization of foreseen events.

Theorem 6 *If conditional preferences $\{\succeq_{t,\omega}\}$ satisfy Order, Consequentialism, Temporal State Independence and Dynamic Consistency, then all events must be foreseen. That is, for all effectively certain acts h in $\mathcal{H}$, it is necessarily true that $h \sim_0 h(\omega)$ for all ω in Ω.*

It seems intuitive that any meaningful definition of foreseen events must be related to the intertemporal consistency of behavior. For example, imagine that you observe conditional preferences at every node and find that behavior is dynamically consistent. If behavior is also consequentialist, the individual necessarily observes and recognizes any event that has transpired. That is, he updates his ranking of actions in response to objective information about the state of the world. The consistency of his behavior then 'reveals' that the individual has foreseen all possible changes of the environment and incorporated them into his plans. Theorem 6 shows that the definition of foreseen events based on the ranking of effectively certain acts is consistent with that intuition.

Theorem 6 and the above discussion focus on the special case when preferences are *fully* dynamically consistent. To extend the analysis to more general cases, it is useful to introduce a notion of *partial* consistency. For any t and ω, let $\overline{\mathcal{G}}^{t,\omega}$ denote a subtree emanating from the node $\mathcal{F}_t(\omega)$, and say that an act h is $\overline{\mathcal{G}}^{t,\omega}$-adapted if its continuation act from the node $\mathcal{F}_t(\omega)$ is $\overline{\mathcal{G}}^{t,\omega}$-adapted.

Definition 7 *Conditional preferences $\{\succeq_{t,\omega}\}$ are **dynamically consistent***

relative to *$\{\overline{\mathcal{G}}^{t,\omega}\}$ if, for all ω and $t < T$, $g \succeq_{t+1,\omega'} g'$ for all $\omega' \in \mathcal{F}_t(\omega)$ implies $g \succeq_{t,\omega} g'$, whenever the acts g, g' are $\overline{\mathcal{G}}^{t,\omega}$-adapted and $g_\tau = g'_\tau$ for all $\tau \leq t$.*

How should partial consistency be interpreted? Throughout the paper, the objective has been to model a sophisticated individual who may not foresee all relevant aspects of the environement but behaves 'rationally' given his perception of the world. Consistent with that view is the interpretation that any violation of time consistency reflects an optimal adjustment in response to some unanticipated change of the environment. Thus, an event A is foreseen in period $t = 0$ only if all bets on A are consistently evaluated over time. Conversely, the inconsistent evaluation of *any* such bet renders the event A unforeseen.

Restating the definition slightly and taking into account that foresight evolves over time, one can define the process of foreseen events as the *largest* process $\{\overline{\mathcal{G}}^{t,\omega}\}$ relative to which behavior is dynamically consistent. Two qualifications turn out to be necessary when searching for the largest process. First, attention must be restricted to the subclass of processes $\{\overline{\mathcal{G}}^{t,\omega}\}$ satisfying the following **regularity** property:[7]

$$\overline{\mathcal{G}}_T^{t,\omega} = \overline{\mathcal{G}}_{T-1}^{t,\omega} \text{ for every } t \text{ and } \omega. \tag{4.1}$$

The necessity of this restriction goes back to the main message of the paper. That is, the empirical implications of unforeseen contingencies *vis-a-vis* ambiguity can be disentangled only in a multi-period setting. More precisely, any behavioral test as to whether an event A is unforeseen or not is powerful only if the event belongs to at least two algebras $\mathcal{F}_t$ and $\mathcal{F}_{t'}$ in the objective tree. In a finite horizon setting, this implies that tests are powerful only if the event belongs to $\mathcal{F}_{T-1}$.

One can be more specific and ask why intertemporal consistency in particular cannot identify whether an event $A \in \mathcal{F}_T \backslash \mathcal{F}_{T-1}$ is foreseen or not. The reason is that, in the last two periods, dynamic consistency obtains trivially in the sense that it reduces to the static property of monotonicity. As argued earlier, a behavioral definition of foreseen events cannot be based on

[7] The assumption that $\mathcal{F}_T = \mathcal{F}_{T-1}$ is not needed for any of the results in this section. If the assumption is nonetheless maintained, then all sequentially connected subtrees satisfy (4.1) by virtue of property (2.1). This motivates the term 'regular'.

the monotonicity of preference. If the individual has zero foresight and evaluates every action by its worst possible outcome, his preferences would still be weakly monotone.

In addition to regularity, one must restrict attention to processes $\{\overline{\mathcal{G}}^{t,\omega}\}$ that are refining in the sense of Definition 5. To understand why the restriction is necessary, consider an environment in which no information is revealed in either of the first two periods, i.e., $\mathcal{F}_0 = \mathcal{F}_1 = \{\Omega, \varnothing\}$. It is then plausible, but not necessary, that conditional preferences in those periods are identical. If that is the case, the subset $\{\succeq_0, \succeq_1\}$ would be trivially dynamically consistent, and yet, its consistency would reveal nothing about the individual's foresight. In fact, an injudicious appeal to Theorem 6 could lead to the potentially erroneous conclusion that the individual has perfect foresight. To narrow down what is foreseen, the entire process of conditional preferences $\{\succeq_{t,\omega}\}$ has to be taken into account. A behavioral distinction between foreseen and unforeseen events is then achieved by looking at those periods in which objective information arrives and behavior, as mandated by Consequentialism, adjusts in response. In terms of the example, if some or all uncertainty resolves in period $t = 2$, one can ignore the uninformative period $t = 0$ and use the subset of preferences $\{\succeq_1, \succeq_{2,\omega}\}$ to identify what is foreseen in period $t = 1$. But since preferences in period $t = 0$ and $t = 1$ are identical, any behavioral definition must single out exactly the same events as being foreseen in period $t = 0$.

More generally, any behavioral definition based on intertemporal consistency must use foresight in later periods as an 'upper bound' in order to narrow down what is foreseen in earlier periods in which intertemporal consistency might have no bite. This is what the restriction to refining processes amounts to.

Under a minimal set of axioms, the corollary below establishes the equivalence of the definitions based on effective certainty and intertemporal consistency respectively. In particular, the equivalence holds for the model of limited foresight developed in Sections 2 and 3.

Corollary 7 *If a family of preferences $\{\succeq_{t,\omega}\}$ satisfies Order, Consequentialism, Temporal State Independence and Weak Dynamic Consistency, then $\{\mathcal{G}^{t,\omega}\}$ is the largest regular refining process relative to which $\{\succeq_{t,\omega}\}$ is dynamically consistent.*

5 Related Literature

In the literature on unforeseen contingencies, a pivotal point concerns the interpretation of the primitive state space Ω. Savage, whose primary interest was normative and whose domain is adopted here, interpreted the state space as a construct used by the individual to make better decisions. As part of a positive theory, however, the state space has to be interpreted as a description of 'objective reality' in the sense that it is constructed by the modeler. The theory developed here should be understood in this sense. Consequently, the question whether the individual knows the state space or not becomes an essentially empirical question which can be tested and answered within the model.

As a way of explicating the point, one may draw a comparison with the analysis in Kreps [13]. He takes as primitive a state space Ω which he interprets as representing the individual's incomplete perception of the environment. To infer if there are other 'unforeseen' contingencies, Kreps asks how much flexibility the individual would be willing to preserve contingent on any of the prespecified states $\omega \in \Omega$. The more flexibility the individual wishes to preserve, the more incomplete a description of reality ω is. Kreps proceeds to derive a Savage-style representation with an extended state space $\Omega \times \Theta$ where the contingencies θ arise endogenously, that is, as part of the representation. The extended state space 'supplies' those details missing from the individual's perception of the environment as captured by Ω. As Kreps [13, p. 278] acknowledges, however, the problem whether the derived contingencies θ are foreseen or not is entirely semantic. In his words, what he calls 'unforeseen', another may well wish to call 'uncontractible'. According to the latter view, 'it isn't that the individual does not foresee the contingencies [θ], but only that we don't allow him to have his consumption so finely conditioned as he would like. So, it can be argued, we don't have a model of unforeseen contingencies at all, but rather a standard model of incomplete contracts.'

Logically, the question in this paper precedes the development in Kreps [13]. That is, the paper asks what behavior would reveal whether the contingencies in the *primitive* state space Ω are foreseen. Notice that the question remains relevant even if one follows Kreps and interprets Ω as a list of contingencies compiled by the individual. From the standpoint of revealed preference theory, any statement about the individual's knowledge or perception should be expressible in terms his choice behavior.

Is there a potential synergy of the two models? To answer the question, suppose some contingencies in Ω turn out to be unforeseen as defined in this paper. Then, at the level of static choice, it would engender no inconsistency to follow Kreps and ask how much flexibility the individual would be willing to preserve contingent on the events in Ω he does foresee. After all, as argued by Kreps, this is the kind of behavior that reveals whether the individual is acting in anticipation of the unforeseen. The approach is not pursued here as it appears to be generally unclear how to incorporate the derived contingencies θ in a theory of dynamic choice. Kreps raises the problem but provides no answer. To be more specific, if θ is not a primitive of the model, it seems that neither can be any preference profile intended to represent behavior 'conditional on θ'. This is a severe limitation in so far as interest in unforeseen contingencies stems primarily from the problem of adaptation. In contrast, the development of a complete model of dynamic choice is made possible here partly because all contingencies, foreseen or not, are part of the primitive state space Ω.

To avoid the interpretational and methodological problems in working with an endogenously extended state space, the paper pursues a different approach to modeling 'anticipation of the unforeseen'. This is essentially the approach developed in Gilboa and Schmeidler [11] and Epstein, Marinacci, and Seo [7]. In these papers, the multiple prior model of ambiguity aversion is interpreted as a 'reduced-form' model of an individual who is aware that there contingencies he does not foresee. The contribution of this paper is to show that the two conceptually different phenomena, ambiguity aversion and limited foresight, are also empirically different while preserving a useful crossover between the respective models.

6 Appendix - Incomplete

All algebras in the appendix are finitely generated. The corresponding simplex Δ of probability measures is endowed with the standard Euclidean topology and Δ° denotes its interior.

A filtration $\{\mathcal{G}_t\}$ on a state space Ω is identified with the algebra $\mathcal{G}$ on $\Omega \times \mathcal{T}$ generated by the sets $A \times \{t\}$ for $A \in \mathcal{G}_t$ and $t \in \mathcal{T}$. Under this identification, an act h is $\{\mathcal{G}_t\}$-adapted if and only if the mapping $(\omega, t) \longmapsto$

$h_t(\omega)$ is $\mathcal{G}$-measurable. The symbol $\mathcal{G}$ is used interchangeably to denote the algebra on $\Omega \times \mathcal{T}$ and the filtration $\{\mathcal{G}_t\}$ on Ω.

For any $A \subset \Omega, x \in M$ and $f \in M^\Omega$, $fAx := f' \in M^\Omega$ where $f'(\omega) = f(\omega)$ if $\omega \in A$, and $f'(\omega) = x$ if $\omega \in A^c$. For any $A \in \mathcal{F}_t, x \in M$ and $h \in \mathcal{H}$, $hAx := (h_{-t}, h_tAx)$.

6.1 Proof of Theorem 2

Adopt the arguments in Epstein and Schneider [5, Lemma A.1] to deduce that $\succeq$ has a representation:

$$U(h) = \min_{q \in Q} \sum_t \beta^t \langle q_t, h_t \rangle . \tag{6.1}$$

Above, Q is a closed, convex subset of $\times_{t \in \mathcal{T}} \Delta(\Omega, \mathcal{F}_t)$. Denote a generic element in Q by $q := (q_t)_{t \in \mathcal{T}}$. For every subset $\mathcal{T}' \subset \mathcal{T}$, $proj_{\mathcal{T}'}(q)$ denotes the vector $(q_t)_{t \in \mathcal{T}'}$. Without loss of generality, set $\beta = 1$ and define $\langle q, h \rangle := \sum_t \langle q_t, h_t \rangle$.

A subfiltration $\mathcal{G}$ of $\mathcal{F}$ defines the following subspace of $\times_{t \in \mathcal{T}} ba(\Omega, \mathcal{F}_t)$:

$$diag(\mathcal{G}) := \{q \in \times_{t \in \mathcal{T}} ba(\Omega, \mathcal{F}_t) : \text{marg}_{\mathcal{G}_t} q_{t+1} = \text{marg}_{\mathcal{G}_t} q_t \text{ for all } t < T\}.$$

Note that $diag(\mathcal{G}) \neq diag(\mathcal{G}')$ whenever $\mathcal{G}_t \neq \mathcal{G}'_t$ for some $t < T$. Equivalently, there exists a bijection between the diagonals of $\times_{t \in \mathcal{T}} ba(\Omega, \mathcal{F}_t)$ and the set of subfiltrations $\mathcal{G}$ such that $\mathcal{G}_T = \mathcal{G}_{T-1}$. Call such subfiltrations *regular*. For a regular subfiltration $\mathcal{G}$, the following lemma establishes a basic duality between the diagonal $diag(\mathcal{G})$ and the set of effectively certain acts in $\mathcal{H}_\mathcal{G}$. In view of (6.1) and after appropriate normalization, the latter can be identified with the subset:

$$\mathcal{H}^c := \{h \in \mathcal{H} : \sum_t h_t(\omega) = 0 \text{ for all } \omega \in \Omega\}.$$

Lemma 8 *$q \in diag(\mathcal{G})$ if and only if $\langle q, h \rangle = 0$ for all $h \in \mathcal{H}_\mathcal{G} \cap \mathcal{H}^c$.*

Proof: To prove necessity observe that for any $h \in \mathcal{H}_\mathcal{G}$ and $q \in diag(\mathcal{G})$,

$$\langle q, h \rangle = \sum_t \langle q_t, h_t \rangle = \sum_t \langle q_T, h_t \rangle = \langle q_T, \sum_t h_t \rangle . \tag{6.2}$$

If $h \in \mathcal{H}_\mathcal{G} \cap \mathcal{H}^c$, then $\sum_t h_t(\omega) = 0$ for all ω and $\langle q, h \rangle = \langle q_T, \sum_t h_t \rangle = 0$ for all $q \in diag(\mathcal{G})$.

To establish sufficiency, first prove that for any $h \in \mathcal{H}_{\mathcal{G}}$

$$h \in \mathcal{H}^c \text{ if and only if } \langle q, h \rangle = 0 \text{ for all } q \in diag(\mathcal{G}). \tag{6.3}$$

Sufficiency of (6.3) follows (6.2). To see the reverse implication, fix some $h \in \mathcal{H}_{\mathcal{G}} \backslash \mathcal{H}^c$ and without loss of generality suppose that $\sum_t h_t(\omega) > 0$ for some ω. Let q_T be a measure in $\Delta(\Omega, \mathcal{G}_T)$ such that $q_T(\mathcal{G}_T(\omega)) = 1$. Since $h \in \mathcal{H}_{\mathcal{G}}$, $\langle q_T, \sum_t h_t \rangle$ is well-defined and strictly positive. Extend q_T to a vector $q \in diag(\mathcal{G})$ and note that $\langle q, h \rangle = \langle q_T, \sum_t h_t \rangle > 0$, proving the necessity of (6.3).

To complete the proof of the lemma, fix some $q' \notin diag(\mathcal{G})$. It suffices to find an act $h' \in \mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c$ such that $\langle q', h' \rangle \neq 0$. Since $diag(\mathcal{G})$ is a subspace of $\times_{t \in \mathcal{T}} ba(\Omega, \mathcal{G}_t)$ and any subspace is the intersection of the hyperplanes that contain it, there exists an act $h' \in \mathcal{H}_{\mathcal{G}}$ such that

$$\langle q', h' \rangle \neq 0 \text{ and } \langle q, h' \rangle = 0 \text{ for all } q \in diag(\mathcal{G}).$$

By (6.3), the act h' lies in $\mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c$, as desired.■

Lemma 9 *For every closed set* $Q \subset \times_{t \in \mathcal{T}} \Delta(\Omega, \mathcal{G}_t)$*,*

$$Q \subset diag(\mathcal{G}) \textit{ if and only if } \min_{q \in Q} \langle q, h \rangle = 0 \textit{ for all } h \in \mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c.$$

Proof: Sufficiency follows directly from Lemma 8. To see necessity, suppose there exists some $q' \in Q \backslash diag(\mathcal{G})$. By Lemma 8, there exists $h' \in \mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c$ such that $\langle q', h' \rangle \neq 0$. Since

$$h' \in \mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c \text{ if and only if } -h' \in \mathcal{H}_{\mathcal{G}} \cap \mathcal{H}^c,$$

one can choose h' such that $\min_{q \in Q} \langle q, h' \rangle \leq \langle q', h' \rangle < 0$. This establishes a contradiction.■

Define the subset of subjectively certain acts:

$$\mathcal{H}^* := \{h \in \mathcal{H}^c : h' \sim h'(\omega) \text{ for all } \omega \in \Omega \text{ and all } h' \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}\},$$

and let $\mathcal{G}^*$ be the algebra on $\Omega \times \mathcal{T}$ induced by $\mathcal{H}^*$.

Lemma 10 *The algebra* $\mathcal{G}^*$ *on* $\Omega \times \mathcal{T}$ *is a regular filtration on* Ω*.*

Proof: First prove that for every $h \in \mathcal{H}^c$, the filtration $\mathcal{F}(h)$ is the smallest algebra on $\Omega \times \mathcal{T}$ induced by $\mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$. For every $t < T$, the act $(\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_\Omega, -\mathbf{1}_\Omega) \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$ implying that the smallest algebra contains the set $\Omega \times \{t\}$ for every $t \in \mathcal{T}$. Also,

$$h \in \mathcal{H}^c \text{ if and only if } h_T = -\sum_{\tau<T} h_\tau.$$

Conclude that for every $h \in \mathcal{H}^c$, $\sigma(h_T) \leq \vee_{\tau \leq T-1}\sigma(h_\tau)$. But then

$$\mathcal{F}(h)_T = \vee_{\tau \leq T}\sigma(h_\tau) = [\vee_{\tau \leq T-1}\sigma(h_\tau)] \vee \sigma(h_T) = \vee_{\tau \leq T-1}\sigma(h_\tau) = \mathcal{F}(h)_{T-1}.$$

Conclude that $\mathcal{F}(h)$ is regular and for all events $A \in \mathcal{F}(h)_T$ and payoffs $x \in M$ (in particular $x \neq 0$):

$$(\mathbf{0}_{-(T-1),-T}, xA^c(-x), xA(-x)) \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)},$$

and, since $\mathcal{F}(h)$ is a filtration, for all $t < T$ and $A \in \mathcal{F}(h)_t$,

$$(\mathbf{0}_{-t,-(t+1)}, xA^c(-x), xA(-x)) \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}.$$

The above inclusions imply that $\mathcal{F}(h)$ is the smallest algebra induced by $\mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$.

Finally, fix $h \in \mathcal{H}^*$, $h' \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$, and $h'' \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h')}$. Since $\mathcal{F}(h') \leq \mathcal{F}(h)$, $h'' \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$. By the choice of h, the latter implies $h'' \sim h''(\omega)$ for all ω. Conclude that $h' \in \mathcal{H}^*$. But then

$$\begin{aligned} \mathcal{H}^* &= \cup_{h \in \mathcal{H}^*}[\mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}] \Rightarrow \\ \mathcal{G}^* &= \vee_{h \in \mathcal{H}^*}\mathcal{F}(h). \end{aligned}$$

Since the supremum of regular filtrations is a regular filtration, the lemma is proved.■

Lemma 11 *$Q \cap diag(\mathcal{G}^*) \neq \varnothing$.*

Proof: Let $Q' = \{(\text{marg}_{\mathcal{G}_t^*} q_t)_{t \in \mathcal{T}} : (q_t)_{t \in \mathcal{T}} \in Q\}$ and define the linear functional

$$\phi : (q_t)_{t>0} \longmapsto (\text{marg}_{\mathcal{G}_{t-1}^*} q_t)_{t>0}.$$

Consider the following subdomains of acts

$$\begin{aligned} D_{\mathcal{T}\setminus\{T\}} &:= \{(h_0, ..., h_{T-1}, x_0) \in \mathcal{H}_{\mathcal{G}^*}\}, \\ D_{\mathcal{T}\setminus\{0\}} &:= \{(x_0, h_0, ..., h_{T-1}) : (h_0, ..., h_{T-1}, x_0) \in \mathcal{H}_{\mathcal{G}^*}\}. \end{aligned}$$

Under the obvious identification, $D_{\mathcal{T}\setminus\{T\}} = D_{\mathcal{T}\setminus\{0\}}$. The restrictions of $\succeq$ to $D_{\mathcal{T}\setminus\{T\}}$ and $D_{\mathcal{T}\setminus\{0\}}$, respectively, are represented by the following utility functions:

$$\begin{aligned} U_{D_{\mathcal{T}\setminus\{T\}}} &=: \min_{q \in proj_{\mathcal{T}\setminus\{T\}} Q'} \textstyle\sum_{t<T} \langle q_t, h_t \rangle \\ U_{D_{\mathcal{T}\setminus\{0\}}} &=: \min_{q \in \phi \circ proj_{\mathcal{T}\setminus\{0\}} Q'} \textstyle\sum_{t>0} \langle q_t, h_t \rangle \end{aligned}$$

By Stationarity, $U_{D_{\mathcal{T}\setminus\{T\}}}$ and $U_{D_{\mathcal{T}\setminus\{0\}}}$ represent the same preference relation. [10, Theorem 1] implies that

$$proj_{\mathcal{T}\setminus\{T\}} Q' = \phi \circ proj_{\mathcal{T}\setminus\{0\}} Q' =: K$$

Define the correspondence

$$\psi := \phi \circ proj_{\mathcal{T}\setminus\{0\}} \circ \left(Q' \cap proj^{-1}_{\mathcal{T}\setminus\{T\}}\right) : K \rightrightarrows K$$

Since Q' is closed, ψ is the composition of a continuous function and an upper hemicontinuous correspondence. Thus ψ is upper hemicontinuous. Since Q' is convex, ψ is also convex-valued. By the Kakutani fixed point theorem [1, Corollary 16.51], ψ has a fixed point $q \in \psi(q)$. Equivalently, there exists a point $(q_0, q_1, ..., q_{T-1}, q_T) \in Q'$ such that

$$\begin{aligned} \phi(q_1, ..., q_T) &= (q_0, q_1, ..., q_{T-1}) \\ &\Leftrightarrow \\ \text{marg}_{\mathcal{G}^*_{t-1}} q_t &= q_{t-1}, \forall t > 1. \blacksquare \end{aligned}$$

Lemma 12 *$\mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*} = \mathcal{H}^*$ and $\mathcal{G}^*$ is the largest regular filtration $\mathcal{G}$ such that $Q \subset diag(\mathcal{G})$.*

Proof: By construction, $\mathcal{H}^* \subset \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$. To see the reverse inclusion, fix $h \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$ and let $x \sim (-h)$. Certainty Independence implies that $\frac{1}{2}x + \frac{1}{2}h \sim \frac{1}{2}(-h) + \frac{1}{2}h$. The two indifferences imply

$$\begin{aligned} x &= -\frac{1}{T+1} \max_{q \in Q} \textstyle\sum_t \langle q_t, h_t \rangle, \\ x &= -\frac{1}{T+1} \min_{q \in Q} \textstyle\sum_t \langle q_t, h_t \rangle. \end{aligned}$$

Conclude that for every $h \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$, $\langle q, h\rangle = \langle q', h\rangle$ for every $q, q' \in Q$. By Lemma 11, there exists a $q \in Q \cap diag(\mathcal{G}^*)$. It follows that $\langle q, h\rangle = 0$ for every $q \in Q$ and so $\mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*} \subset \mathcal{H}^*$. Also by Lemma 9, $Q \subset diag(\mathcal{G}^*)$.

If $\mathcal{G}$ is any filtration such that $\mathcal{G}_T = \mathcal{G}_{T-1}$ and $Q \subset diag(\mathcal{G})$, Lemma 9 implies that

$$\mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}} \subset \mathcal{H}^* = \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$$

Conclude that $\mathcal{G} \subset \mathcal{G}^*$.■

6.1.1 Properties of the Filtration $\mathcal{G}^*$

Lemma 13 *$\mathcal{G}^*$ is connected, that is, $\mathcal{G}^*_t = \mathcal{G}^*_T \cap \mathcal{F}_t$ for all $t \in \mathcal{T}$.*

Proof: By Lemma 10, $\mathcal{G}^*$ is a filtration. Thus, $\mathcal{G}^*_t \subset \mathcal{G}^*_T \cap \mathcal{F}_t$ for all t. Conversely, fix an event $A \in \mathcal{G}^*_T \cap \mathcal{F}_t$ for some $t \in \mathcal{T}$. Since $\mathcal{G}^*$ is regular by Lemma 10, $\mathcal{G}^*_{T-1} = \mathcal{G}^*_T \ni A$ for $t = T-1$. For $t < T-1$, it suffices to show that $(\mathbf{0}_{-t,-(t+1)}, xAy, (-x)A(-y)) \sim \mathbf{0}$ for all $x, y \in M$. By the regularity of $\mathcal{G}^*$, $(\mathbf{0}_{-(T-1),-T}, xAy, (-x)A(-y)) \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$, which implies that

$$(\mathbf{0}_{-(T-1),-T}, xAy, (-x)A(-y)) \sim \mathbf{0}.$$

Applying Stationarity repeatedly, conclude that

$$(\mathbf{0}_{-t,-(t+1)}, xAy, (-x)A(-y)) \sim \mathbf{0}.■$$

Lemma 14 *$\mathcal{G}^*$ is sequentially connected.*

Proof: Fix an event $A \in \Pi_{\mathcal{G}_t} \backslash \Pi_{\mathcal{F}_t}$ for some $t < T$. By way of contradiction, suppose there exists a set $\varnothing \neq B \in \mathcal{G}_{t+1}$ such that $B \subsetneq A$. First, suppose $B \subset C \subsetneq A$ for some $C \in \Pi_{\mathcal{F}_t}$. Since $\mathcal{G}$ is connected, conclude that $B \subsetneq C$. Otherwise, $B = C \in \mathcal{G}_{t+1} \cap \Pi_{\mathcal{F}_t}$ implies that $C \in \Pi_{\mathcal{G}_t}$, contradicting the choice of A. Now take the acts $g = (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_B)$ and $g' = (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_C)$. By construction, $g \in \mathcal{H}_{\mathcal{G}}$ and g' simplifies g at $C \in \Pi_{\mathcal{F}_t}$. By Simplification, $g' \in \mathcal{H}_{\mathcal{G}}$ and so $C \in \mathcal{G}_{t+1} \cap \Pi_{\mathcal{F}_t}$. Since $\mathcal{G}$ is connected, $C \in \Pi_{\mathcal{G}_t}$ contradicting $C \subsetneq A \in \Pi_{\mathcal{G}_t}$.

Conversely, suppose $B \cap C \neq \varnothing$ and $B \cap C^c \neq \varnothing$ for some $C \in \Pi_{\mathcal{F}_t}$. The act $g = (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_B)$ is $\mathcal{G}^*$-measurable. Moreover, the continuation of $g' := (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_{B \cup C})$ at the node $C \in \Pi_{\mathcal{F}_t}$ simplifies the continuation of g . By Simplification, $B \cup C \in \mathcal{G}_{t+1}$. But $B \cup C \in \mathcal{G}_{t+1}$ and $B \in \mathcal{G}_{t+1}$ imply $C \backslash B = B \cup C \backslash B \in \mathcal{G}_{t+1}$ and, by construction, $\varnothing \neq C \backslash B \subsetneq C \subsetneq A$ and $C \in \Pi_{\mathcal{F}_t}$. But then $g' = (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_C)$ simplifies $g = (\mathbf{0}_{-t,-(t+1)}, \mathbf{1}_A, \mathbf{1}_{C \backslash B}) \in \mathcal{H}_{\mathcal{G}}$ at $C \in \Pi_{\mathcal{F}_t}$. As before, conclude that $C \in \Pi_{\mathcal{G}_t}$ contradicting $C \subsetneq A \in \Pi_{\mathcal{G}_t}$.■

6.1.2 Construction of the Approximation Mapping Φ

Lemma 15 *For all acts h, h', payoffs x and y and events $A \in \cup_t \Pi_{\mathcal{G}_t^*}$,*

$$hAx \succeq h'Ax \text{ if and only if } hAy \succeq h'Ay$$

Proof: Suppose $hAx \succeq h'Ax$ and note that $(hAy)Ax = hAx$ and $(h'Ay)Ax = h'Ax$. Conclude that $(hAy)Ax \succeq (h'Ay)Ax$. For any $A' \in \cup_t \Pi_{\mathcal{G}_t^*}$ and $A' \neq A$,

$$(hAy)A'x = (h'Ay)A'x = yA'x.$$

Conclude that

$$(hAy)A'x \sim (h'Ay)A'x.$$

Thus, $(hAy)A'x \succeq (h'Ay)A'x$ for all $A' \in \cup_t \Pi_{\mathcal{G}_t^*}$. By Subjective Monotonicity, $hAy \succeq h'Ay$ as desired.■

Fix some x_0 and for every $A \in \cup_t \Pi_{\mathcal{G}_t^*}$, define the preference $\succeq_A$ as

$$h \succeq_A h' \text{ if and only if } hAx_0 \succeq h'Ax_0.$$

By the above lemma, the 'conditional' preference $\succeq_A$ is independent of the choice of x_0. By construction, $\succeq_A$ inherits convexity, monotonicity and mixture-continuity. By Nonnullity, the preference is also nontrivial. By Certainty Independence and by Lemma 15 in turn,

$$\begin{aligned} hAx_0 &\succeq h'Ax_0 \Rightarrow \\ [\alpha h + (1-\alpha)x]A[\alpha x_0 + (1-\alpha)x] &\succeq [\alpha h' + (1-\alpha)x]A[\alpha x_0 + (1-\alpha)x] \Rightarrow \\ [\alpha h + (1-\alpha)x]Ax_0 &\succeq [\alpha h' + (1-\alpha)x]Ax_0 \end{aligned}$$

Conclude that for all $A \in \cup_t \Pi_{\mathcal{G}_t^*}$, $\succeq_A$ is a multiple prior preference. Next, recall that, for any (ω, t) in $\Omega \times \mathcal{T}$, $\mathcal{G}_t^*(\omega)$ denotes the event in $\Pi_{\mathcal{G}_t^*}$ containing ω. By [10, Theorem 1], there exists a set $\mathcal{C}_{\omega,t} \subset \Delta(\mathcal{G}_t^*(\omega), \mathcal{F}_t)$ such that

$$h \succeq_{\mathcal{G}_t^*(\omega)} h' \text{ if and only if } \min_{q \in \mathcal{C}_{\omega,t}} \langle q, h_t \rangle \geq \min_{q \in \mathcal{C}_{\omega,t}} \langle q, h_t' \rangle . \quad (6.4)$$

Define the mapping $\Phi : \mathcal{H} \to \mathcal{H}$

$$(\Phi h)_t(\omega) = \min_{q \in \mathcal{C}_{\omega,t}} \langle q, h_t \rangle .$$

By construction,

$$\Phi(hAx_0) = \Phi(h)Ax_0 \sim hAx_0 \text{ for all } A \in \cup_t \Pi_{\mathcal{G}_t^*}.$$

By Subjective Monotonicity, $\Phi(h) \sim h$ for all $h \in \mathcal{H}$. By Lemma 1, Φ is an approximation mapping. To conclude the proof of the theorem, for all $h \in \mathcal{H}$, define

$$\begin{aligned} V(h) &= U \circ \Phi(h) \\ &= \min_{q \in Q} \textstyle\sum_t \langle q_t, \Phi(h)_t \rangle \\ &= \min_{q \in Q} \textstyle\sum_t \langle q_T, \Phi(h)_t \rangle \\ &= \min_{q \in Q} \langle q_T, \textstyle\sum_t \Phi(h)_t \rangle . \end{aligned}$$

The third equality follows from Lemma 12. Finally set

$$\mathcal{C} := \text{marg}_{\mathcal{G}_T^*} \circ proj_{\{T\}} Q.$$

The claim that $\mathcal{C}$ is a subset of $\Delta^\circ(\Omega, \mathcal{G}_T^*)$ follows from the following property of multiple-prior preferences.

Lemma 16 *Let $\succeq$ be a multiple-prior preference on the space of $\mathcal{F}_T$-measurable functions in M^Ω, and let $\mathcal{C}$ be the respective set of priors. Let Π be any partition such that $\Pi \leq \Pi_{\mathcal{F}_T}$ and suppose that for all payoffs $x \in M$:*

$$hAx \succeq h'Ax \text{ for all } A \in \Pi \text{ implies } h \succeq h'$$

If $A \in \Pi$ is nonnull, then $p(A) > 0$ for all $p \in \mathcal{C}$.

Proof: Suppose by way of contradiction that $p(A) = 0$ for some $A \in \Pi$ and $p \in \mathcal{C}$. Since A is nonnull, $\max_{q \in \mathcal{C}} q(A) > p(A) = 0$. Fix some payoffs y, y' such that $1 > y > y' > 0$ and note that:

$$U(yA0) = \min_{q \in C}[q(A)y] = y \min_{q \in C} q(A) = 0 = U(y'A0), \text{ and}$$

$$U(yA1) = \min_{q \in C}[(y-1)q(A) + 1] = 1 - (1-y)\max_{q \in C} q(A) >$$

$$> 1 - (1-y')\max_{q \in C} q(A) = U(y'A1).$$

Conclude that $yA0 \sim y'A0$ and $yA1 \succ y'A1$ in contradiction of Lemma 15.■

6.1.3 Uniqueness

Uniqueness of the set $\mathcal{C}$ follows from familiar arguments. To prove the uniqueness of the $\mathcal{G}^*$-approximation mapping, take two such mappings Φ, $\widehat{\Phi}$. By separability, for all $A \in \Pi_{\mathcal{G}^*}$, h' and x:

$$\begin{aligned}
\widehat{\Phi}(h'Ax) &= \widehat{\Phi}((h'Ax)Ax)A\widehat{\Phi}((h'Ax)A^c x) \\
&= \widehat{\Phi}(h'Ax)A\widehat{\Phi}(x) \\
&= \widehat{\Phi}(h'Ax)Ax
\end{aligned}$$

The last equality follows from the fact that $\widehat{\Phi}$ must be identity on $\mathcal{G}^*$-measurable acts. Since $A \in \cup_t \Pi_{\mathcal{G}^*_t}$ and $\Phi(h'Ax)$, $\widehat{\Phi}(h'Ax) \in \mathcal{H}_{\mathcal{G}^*}$, there exist payoffs y, $\widehat{y}$ such that

$$\begin{aligned}
\Phi(h'Ax) &= yAx \\
\widehat{\Phi}(h'Ax) &= \widehat{y}Ax
\end{aligned}$$

Since $\widehat{\Phi}(h'Ax) \sim \Phi(h'Ax)$ and $\succeq$ is strictly increasing on $\mathcal{H}_{\mathcal{G}^*}$, it must be the case that $y = \widehat{y}$. The proof is completed by induction on the number of events $A \in \Pi_{\mathcal{G}^*}$ such that an act $h' \in \mathcal{H}$ is nonconstant.

6.2 An Alternative Formulation

This section describes an alternative formulation of the static model.

Definition 8 *A preference relation $\succeq$ on $\mathcal{H}$ has a **regular representation** $(\mathcal{G}, \Phi, \mathcal{C})$ if it admits a utility function V of the form (2.3) where $\mathcal{G}$ is regular, the mapping Φ is identity on $\mathcal{H}_{\mathcal{G}}$ and $\mathcal{C}$ is a closed, convex subset of $\Delta^\circ(\Omega, \mathcal{G}_T)$.*

Definition 9 *A preference relation $\succeq$ on $\mathcal{H}$ has a **largest representation** $(\mathcal{G}, \Phi, \mathcal{C})$ if it admits a utility function V of the form (2.3) where $\mathcal{G}$ is sequentially connected, Φ is a $\mathcal{G}$-approximation mapping, $\mathcal{C}$ is a closed, convex subset of $\Delta^\circ(\Omega, \mathcal{G}_T)$ and $\mathcal{G}$ is the largest filtration for which a regular representation exists.*

Lemma 17 *A preference $\succeq$ has a limited foresight representation if and only if it has a largest representation. The two representations are identical.*

Proof: If $(\mathcal{G}, \Phi, \mathcal{C})$ is a regular representation for $\succeq$, then $\mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}} \subset \mathcal{H}^*$ and so $\mathcal{G} \subset \mathcal{G}^*$. Thus, the limited foresight model is sufficient for a largest representation. Conversely, if a largest representation $(\mathcal{G}^*, \Phi, \mathcal{C})$ exists, then

$$\begin{aligned} \mathcal{G}^* &= \vee_{Q \subset diag(\mathcal{G}')} \mathcal{G}' \Rightarrow \\ \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*} &= \cup_{\{\mathcal{G}' : Q \subset diag(\mathcal{G}')\}} \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}'} \end{aligned}$$

At the same time,

$$\mathcal{H}^* = \cup_{\{\mathcal{G}' : Q \subset diag(\mathcal{G}')\}} \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}'}.$$

To see this note that, $\cup_{\{\mathcal{G}' : Q \subset diag(\mathcal{G}')\}} \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}'} \subset \mathcal{H}^*$. If $h \in \mathcal{H}^*$, then for all $h' \in \mathcal{H}^c \cap \mathcal{H}_{\mathcal{F}(h)}$, $h' \sim \mathbf{0}$. From Lemma 9, conclude that $Q \subset diag(\mathcal{F}(h))$ and so $h \in \cup_{\mathcal{G}' : Q \subset diag(\mathcal{G}')} \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}'}$. Thus $\mathcal{H}^* = \mathcal{H}^c \cap \mathcal{H}_{\mathcal{G}^*}$ and so $\mathcal{G}^* = \mathcal{F}(\mathcal{H}^*)$ as desired.■

6.3 Proof of Lemma 3

Suppose $\succeq$ has a representation $(\mathcal{G}, \{\mathcal{C}_A\}, \mathcal{C})$ such that $\mathcal{C}_A$ has nonempty interior in $\Delta(A, \mathcal{F}_t)$ for all $t \in \mathcal{T}$ and $A \in \Pi_{\mathcal{G}_t}$. Let Q be the set in $\times_t \Delta(\Omega, \mathcal{F}_t)$ that represents $\succeq$ as in (6.1). It suffices to show that for all filtrations $\mathcal{G}'$ such that $\mathcal{G}'_{T-1} = \mathcal{G}'_T$:

$$Q \subset diag(\mathcal{G}') \text{ implies } \mathcal{G}' \leq \mathcal{G}.$$

Equivalently,

$$\forall t < T, \forall B \notin \mathcal{G}_t, \ \exists q \in Q \text{ such that } q_t(B) \neq q_{t+1}(B).$$

For all $t \in \mathcal{T}$ and $A \in \Pi_{\mathcal{G}_t}$, define $\succeq_A$ as in (6.4). The family of preferences $\{\succeq_A\}$ and $\succeq$ satisfy the conditions of [5, Theorem 3.2]. Conclude that:

$$Q = \cup_{\mu \in \mathcal{C}} \{(q_t) : q_t = \int p_A d\mu \text{ for some selection } \{p_A\}_{A \in \Pi_{\mathcal{G}_t}} \text{ s.t. } p_A \in \mathcal{C}_A\}.$$

Fix $t < T$ and $B \notin \mathcal{G}_t$ and any $\mu \in \mathcal{C}$. By the above decomposition of Q, it suffices to find two selections $\{p_A\}_{A \in \Pi_{\mathcal{G}_t}}$ and $\{p'_A\}_{A \in \Pi_{\mathcal{G}_t}}$ such that:

$$\int p_A(B) d\mu \neq \int p'_A(B) d\mu.$$

Since $B \notin \mathcal{G}_t$, there exists $A^* \in \Pi_{\mathcal{G}_t}$ such that $A^* \neq B \cap A^* \neq \varnothing$. Since $\mathcal{C}_{A^*}$ has nonempty interior, there exist p_{A^*} and p'_{A^*} such that $p_{A^*}(B) \neq p'_{A^*}(B)$. Choose any $p_A = p'_A$ for all $A \neq A^*$ and $A \in \Pi_{\mathcal{G}_t}$ to complete the proof of the lemma.

6.4 Proof of Theorem 4

Necessity is standard. To prove sufficiency, first show that $\{\mathcal{G}^{t,\omega}\}$ is refining. Fix t, ω such that $\mathcal{F}_{t+1}(\omega) \notin \mathcal{G}^{t,\omega}$. Since $\mathcal{G}^{t,\omega}$ is sequentially connected:

$$\{B \cap \mathcal{F}_{t+1}(\omega) : B \in \mathcal{G}^{t,\omega}\} = \{\varnothing, \mathcal{F}_{t+1}(\omega)\}.$$

Conclude that $\mathcal{G}^{t+1,\omega}$ refines the trivial filtration $\mathcal{G}^{t,\omega} \cap \mathcal{F}_{t+1}(\omega)$. Alternatively, suppose $\mathcal{F}_{t+1}(\omega) \in \mathcal{G}^{t,\omega}$ and take an effectively certain act h such that $h \sim_{t,\omega} \mathbf{0}$. One can assume that $h_{t'}(\omega') = 0$ whenever $t' < t$ or $\omega' \notin \mathcal{F}_{t+1}(\omega)$. It suffices to show that $h \sim_{t+1,\omega} \mathbf{0}$. Let $x(\omega)$ be an outcome such that

$$(\mathbf{0}_{-(t+1)}, x(\omega)) \sim_{t+1,\omega} h,$$

and let f be a function on Ω that pays $x(\omega)$ if $\omega' \in \mathcal{F}_{t+1}(\omega)$ and zero else. By Consequentialism,

$$(\mathbf{0}_{-(t+1)}, f) \sim_{t+1,\omega'} h.$$

for every ω'. By construction, $(\mathbf{0}_{-(t+1)}, f)$ is $\mathcal{G}^{t,\omega}$-adapted, and hence, by Weak Dynamic Consistency,

$$(\mathbf{0}_{-(t+1)}, f) \sim_{t,\omega} h \sim_{t,\omega} \mathbf{0}.$$

By Lemma 16, this is possible if and only if $x(\omega) = 0$. But $x(\omega)$ was chosen such that

$$(\mathbf{0}_{-(t+1)}, f) \sim_{t+1,\omega} h$$

Conclude that $h \sim_{t+1,\omega} \mathbf{0}$, as desired.

To prove that $\{\mathcal{C}^{t,\omega}\}$ admits a consistent extension, construct the set $\mathcal{C}$ recursively. For all ω and $t \geq T-1$, set $\widehat{\mathcal{C}}^{t,\omega} := \mathcal{C}^{t,\omega}$. Fix ω and $t < T-1$ and suppose $\widehat{\mathcal{C}}^{t+1,\omega'}$ has been defined for all ω'. For any ω' such that $\mathcal{F}_{t+1}(\omega') \notin \mathcal{G}^{t,\omega}$, fix some measure $\lambda_{\omega'} \in \Delta^{\circ}(\mathcal{G}^{t,\omega}(\omega'), \mathcal{F}_{t+1})$ such that $\lambda_{\omega''} = \lambda_{\omega'}$ for all $\omega'' \in \mathcal{G}^{t,\omega}(\omega')$. For each $\mu \in \mathcal{C}^{t,\omega}$ define the measure $\widehat{\mu} := \int_{\Omega} \widehat{\lambda}_{\omega'} dm$ in $\Delta^{\circ}(\mathcal{F}_t(\omega), \mathcal{F}_{t+1})$ where

$$m := \operatorname{marg}_{\mathcal{G}^{t,\omega}_{t+1}} \mu \text{ and } \widehat{\lambda}_{\omega'} := \left\{ \begin{array}{l} \lambda_{\omega'} \text{ if } \mathcal{F}_{t+1}(\omega') \notin \mathcal{G}^{t,\omega} \\ \delta_{\omega'} \text{ if } \mathcal{F}_{t+1}(\omega') \in \mathcal{G}^{t,\omega} \end{array} \right. . \tag{6.5}$$

In effect, the constructed measures $\widehat{\mu}$ extend the individual's one-step ahead beliefs at t, ω from the foreseen events in $\mathcal{G}^{t,\omega}_{t+1}$ to all $\mathcal{F}_{t+1}$-measurable subsets of $\mathcal{F}_t(\omega)$. The construction ensures that the set of extensions $M^{t,\omega} := \{\widehat{\mu} : \mu \in \mathcal{C}^{t,\omega}\} \subset \Delta^\circ(\mathcal{F}_t(\omega), \mathcal{F}_{t+1})$ is closed and convex.

Next, let p denote a generic, $\mathcal{F}_{t+1}$-measurable selection from $\omega' \longmapsto \widehat{\mathcal{C}}^{t+1,\omega'}$ and define

$$\widehat{\mathcal{C}}^{t,\omega} = \{\int_\Omega p_{\omega'} d\widehat{\mu}(\omega') : \widehat{\mu} \in M^{t,\omega} \text{ and } p_{\omega'} \in \ \widehat{\mathcal{C}}^{t+1,\omega'} \text{ for all } \omega'\}. \tag{6.6}$$

From [5, Theorem 3.2], conclude that $\widehat{\mathcal{C}}^{t,\omega}$ is a closed and convex subset of $\Delta^\circ(\mathcal{F}_t(\omega), \mathcal{F}_T)$ and $\mathcal{C} := \widehat{\mathcal{C}}^0$ is $\{\mathcal{F}_t\}$-rectangular. In particular,

$$\{\mu(\cdot \mid \mathcal{F}_t(\omega)) : \mu \in \mathcal{C}\} = \widehat{\mathcal{C}}^{t,\omega} \text{ for all } t \text{ and } \omega.$$

To complete the proof, it remains to show that

$$\text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega} := \{\text{marg}_{\mathcal{G}^{t,\omega}} \mu : \mu \in \widehat{\mathcal{C}}^{t,\omega}\} = \mathcal{C}^{t,\omega} \text{ for all } t \text{ and } \omega. \tag{6.7}$$

The next lemmas show that both $\text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega}$ and $\mathcal{C}^{t,\omega}$ admit decompositions similar to (6.6).

Lemma 18 *For all t and ω, $\text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega}$ admits the decomposition*

$$\text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega} = \{\int_\Omega p_{\omega'} d\mu : \mu \in \text{marg}_{\mathcal{G}^{t,\omega}_{t+1}} M^{t,\omega} \textit{ and } p_{\omega'} \in \text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t+1,\omega'}\}.$$

Proof: By (6.6), all measures in $\widehat{\mathcal{C}}^{t,\omega}$ are of the form

$$\int_\Omega p_{\omega'} d\widehat{\mu},$$

where $\widehat{\mu} \in M^{t,\omega}$ and p is an $\mathcal{F}_{t+1}$-measurable selection from the correspondence $\omega' \longmapsto \widehat{\mathcal{C}}^{t+1,\omega'}$. Since $\mathcal{G}^{t,\omega}$ is sequentially connected, $\text{marg}_{\mathcal{G}^{t,\omega}} p$ is a $\mathcal{G}^{t,\omega}_{t+1}$-measurable selection from $\omega' \longmapsto \text{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t+1,\omega'}$. Conclude that

$$\text{marg}_{\mathcal{G}^{t,\omega}} \int_\Omega p_{\omega'} d\widehat{\mu} = \int_\Omega \text{marg}_{\mathcal{G}^{t,\omega}} p_{\omega'} \ d\left(\text{marg}_{\mathcal{G}^{t,\omega}_{t+1}} \widehat{\mu}\right). \blacksquare$$

Lemma 19 *For all t and ω, $\mathcal{C}^{t,\omega}$ admits the decomposition*

$$\mathcal{C}^{t,\omega} = \{\int p_{\omega'} dm : m \in \text{marg}_{\mathcal{G}^{t,\omega}_{t+1}} \mathcal{C}^{t,\omega} \textit{ and } p_{\omega'} \in \text{marg}_{\mathcal{G}^{t,\omega}} \mathcal{C}^{t+1,\omega'} \textit{ for all } \omega'\}.$$

Proof: For each $\omega' \in \mathcal{F}_t(\omega)$, let $\succeq^a_{t+1,\omega'}$ and $\succeq^a_{t,\omega}$ denote the respective restrictions of $\succeq_{t+1,\omega'}$ and $\succeq_{t,\omega}$ to $\mathcal{H}_{\mathcal{G}^{t,\omega}}$. Since $\mathcal{G}^{t+1,\omega'}$ refines $\mathcal{G}^{t,\omega}$ for each $\omega' \in \mathcal{F}_t(\omega)$, the corresponding preference $\succeq^a_{t+1,\omega'}$ admits a representation

$$U^{t+1,\omega'}(h) = \min_{\mu \in \mathrm{marg}_{\mathcal{G}^{t,\omega}} \mathcal{C}^{t+1,\omega'}} \int \left(\sum_{\tau \geq t+1} \beta^{\tau-(t+1)} h_\tau \right) d\mu$$

Since $\mathcal{G}^{t,\omega}$ is sequentially connected, the mapping $\omega' \longmapsto \succeq^a_{t+1,\omega'}$ is $\mathcal{G}^{t,\omega}_{t+1}$-measurable. Thus, the collection of preferences

$$\{\succeq^a_{t,\omega}, \succeq^a_{t+1,\omega'} \colon \omega' \in \mathcal{F}_t(\omega)\}$$

satisfies Consequentialism with respect to the filtration $\mathcal{G}^{t,\omega}$. By State Independence and Lemma 15, the collection of preferences is also dynamically consistent. The claim of the lemma follows from [5, Theorem 3.2].■

Complete the proof of (6.7) by induction. The claim holds trivially for ω and $t \geq T-1$. Fix some ω and $t < T-1$ and suppose the claim has been established for $t+1$. Applying Lemma 19, the induction hypothesis and Lemma 18 in turn, conclude that

$$\begin{aligned} \{\mu(\cdot \mid \mathcal{F}_{t+1}(\omega')) : \mu \in \mathcal{C}^{t,\omega}\} &= \mathrm{marg}_{\mathcal{G}^{t,\omega}} \mathcal{C}^{t+1,\omega'} \\ &= \mathrm{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t+1,\omega'} \\ &= \mathrm{marg}_{\mathcal{G}^{t,\omega}} \{\mu(\cdot \mid \mathcal{F}_{t+1}(\omega')) : \mu \in \widehat{\mathcal{C}}^{t,\omega}\} \end{aligned} \tag{6.8}$$

Also, by construction,

$$\mathrm{marg}_{\mathcal{G}^{t,\omega}_{t+1}} \mathcal{C}^{t,\omega} = \mathrm{marg}_{\mathcal{G}^{t,\omega}_{t+1}} M^{t,\omega}. \tag{6.9}$$

Properties (6.8) and (6.9) show that $\mathrm{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega}$ and $\mathcal{C}^{t,\omega}$ induce the same sets of conditionals and one-step-ahead marginals. By Lemmas 18 and 19, both $\mathrm{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega}$ and $\mathcal{C}^{t,\omega}$ are uniquely determined by the respective sets of conditionals and marginals. Conclude that $\mathrm{marg}_{\mathcal{G}^{t,\omega}} \widehat{\mathcal{C}}^{t,\omega} = \mathcal{C}^{t,\omega}$.

6.4.1 Uniqueness

Let $\mathcal{C}$ be an $\{\mathcal{F}_t\}$-rectangular subset of $\Delta^\circ(\Omega, \mathcal{F}_T)$ and for each t, ω, let $\mathcal{G}^{t,\omega}$ be a sequentially connected algebra such that $\mathcal{F}_{t+1}(\omega') \in \mathcal{G}^{t,\omega}$ for all $\omega' \in \mathcal{F}_t(\omega)$.

Lemma 20 *A measure* μ *in* $\Delta^\circ(\Omega, \mathcal{F}_T)$ *belongs to* $\mathcal{C}$ *if and only if*

$$\mathrm{marg}_{\mathcal{G}^{t,\omega}}\mu(\cdot \mid \mathcal{F}_t(\omega)) \in \{\mathrm{marg}_{\mathcal{G}^{t,\omega}}\mu'(\cdot \mid \mathcal{F}_t(\omega)) : \mu' \in \mathcal{C}\} \textit{ for all } t \textit{ and } \omega.$$

Proof: Sufficiency is immediate. To prove necessity, note the recursive construction of the $\{\mathcal{F}_t\}$-rectangular set in the proof of Theorem 4. An $\{\mathcal{F}_t\}$-rectangular subset contains the measure $\widehat{\mu}$ if and only if

$$\widehat{\mu}(\cdot \mid \mathcal{F}_T(\omega)) \in \{\mu'(\cdot \mid \mathcal{F}_T(\omega)) : \mu' \in \mathcal{C}\} \text{ for } \omega \in \Omega, \text{ and}$$

$$\widehat{\mu}(\cdot \mid \mathcal{F}_t(\omega)) \in \{\mathrm{marg}_{\mathcal{F}_{t+1}}\mu'(\cdot \mid \mathcal{F}_t(\omega)) : \mu' \in \mathcal{C}\} \text{ for } \omega \in \Omega \text{ and } t < T.$$

By construction, the restriction of $\mathcal{F}_T$ to $\mathcal{F}_T(\omega)$ equals $\{\mathcal{F}_T(\omega), \varnothing\}$ which equals $\mathcal{G}^{T,\omega}$ for each $\omega \in \Omega$. By hypothesis, the restriction of $\mathcal{F}_{t+1}$ to $\mathcal{F}_t(\omega)$ is refined by $\mathcal{G}^{t,\omega}$ for each t and ω. Conclude that $\mu \in \mathcal{C}$.■

To prove Theorem 5, let $\{\mathcal{C}^{t,\omega}\}$ be an $\mathcal{F}$-adapted process where $\mathcal{C}^{t,\omega} \subset \Delta^\circ(\mathcal{F}_t(\omega), \mathcal{G}^{t,\omega})$. Say that the set $\mathcal{C}'$ in $\Delta(\Omega, \mathcal{F}_T)$ *extends* $\{\mathcal{C}^{t,\omega}\}$ if

$$\{\mathrm{marg}_{\mathcal{G}^{t,\omega}}\mu(\cdot \mid \mathcal{F}_t(\omega)) : \mu \in \mathcal{C}' \text{ s.t. } \mu(\mathcal{F}_t(\omega)) > 0\} = \mathcal{C}^{t,\omega}.$$

It is not difficult to see that any extension $\mathcal{C}'$ must be a subset of $\Delta^\circ(\Omega, \mathcal{F}_T)$. By way of contradiction suppose that there exists a measure $\mu' \in \mathcal{C}'$ such that $\mu'(\mathcal{F}_t(\omega)) = 0$ for some t and ω. Since for all ω', $\mu'(\mathcal{F}_0(\omega')) = \mu'(\Omega) = 1$, conclude that $t > 0$. Let t^* be the largest t' such that $\mu'(\mathcal{F}_{t^*}(\omega)) > 0$. The time t^* exists since $\mu'(\mathcal{F}_0(\omega)) > 0$. By the definition of t^*, $\mu'(\cdot \mid \mathcal{F}_{t^*}(\omega))$ is well-defined and $\mu'(\mathcal{F}_{t^*+1}(\omega) \mid \mathcal{F}_{t^*}(\omega)) = 0$. The latter gives a contradiction, since $\mathcal{F}_{t^*+1}(\omega) \in \mathcal{G}^{t^*,\omega}$ and

$$\mathrm{marg}_{\mathcal{G}^{t^*,\omega}}\mu'(\cdot \mid \mathcal{F}_{t^*}(\omega)) \in \mathcal{C}^{t^*,\omega} \subset \Delta^\circ(\mathcal{F}_{t^*}(\omega), \mathcal{G}^{t^*,\omega}).$$

From the recursive construction in the proof of Theorem 4, conclude that there exists an $\{\mathcal{F}_t\}$-rectangular extension $\mathcal{C}$ in $\Delta^\circ(\Omega, \mathcal{F}_T)$. Lemma 20 proves that $\mathcal{C}$ is the unique $\{\mathcal{F}_t\}$-rectangular extension, whenever $\mathcal{F}_{t+1}(\omega') \in \mathcal{G}^{t,\omega}$ for all $\omega' \in \mathcal{F}_t(\omega)$ and all t and ω. Moreover, any other, possibly non-rectangular, extension $\mathcal{C}'$ is a subset of $\mathcal{C}$.

6.5 Proof of Theorem 6

To prove Theorem 6, take an act h and suppose $h(\omega) \sim_0 h(\omega')$ for all $\omega, \omega' \in \Omega$. By Consequentialism,

$$h \sim_{T,\omega} h(\omega) \text{ for all } \omega \in \Omega. \tag{6.10}$$

Fix some ω and note that:

$$h_\tau(\omega') = h_\tau(\omega'') \text{ for all } \omega', \omega'' \in \mathcal{F}_{T-1}(\omega) \text{ and } \tau \leq T-1.$$

By State Independence,

$$h(\omega') \sim_{T,\omega} h(\omega'') \text{ for all } \omega', \omega'' \in \mathcal{F}_{T-1}(\omega). \tag{6.11}$$

But then (6.10) and (6.11) imply that for all $\omega' \in \mathcal{F}_{T-1}(\omega)$:

$$h \sim_{T,\omega'} h(\omega') \sim_{T,\omega'} h(\omega).$$

By Dynamic Consistency, $h \sim_{T-1,\omega} h(\omega)$. Proceeding inductively, conclude that $h \sim_0 h(\omega)$.

6.6 Sequentially Connected Filtrations

Say that a filtration $\{\mathcal{G}_t\}$ is connected if

$$\mathcal{G}_t = \mathcal{F}_t \cap \mathcal{G}_T \text{ for all } t \in \mathcal{T}.$$

Proposition 21 *A sequentially connected filtration $\{\mathcal{G}_t\}$ is connected.*

Let $\{\mathcal{G}_t\}$ be sequentially connected. It is evident that $\mathcal{G}_t \subset \mathcal{G}_T \cap \mathcal{F}_t$ for all $t \in \mathcal{T}$. To prove the opposite inclusion, take an event $A \in \mathcal{G}_T \cap \mathcal{F}_t$. If $A \notin \mathcal{G}_t$, then there exists a set $B \in \Pi_{\mathcal{G}_t}$ such that $\varnothing \neq B \cap A \neq B$. Then $A \in \mathcal{F}_t$ implies $B \notin \Pi_{\mathcal{F}_t}$. Conclude that $B \in \Pi_{\mathcal{G}_t} \backslash \Pi_{\mathcal{F}_t}$ and since $\{\mathcal{G}_t\}$ is connected, $B \in \Pi_{\mathcal{G}_T}$. But then $\varnothing \neq B \cap A \neq B$ contradicts the fact that $A \in \mathcal{G}_T$.■

The following proposition shows that sequentially connected filtrations inherit the lattice properties of stopping-times.

Proposition 22 *The class of sequentially connected filtrations is a lattice. It is lattice-isomorphic to the class of sequentially connected algebras.*

First, establish the following distributive law.

Lemma 23 *If $\mathcal{G}$ and $\mathcal{G}'$ are sequentially connected algebras, then*

$$\Pi_{\mathcal{G}\cap\mathcal{F}_t} \vee \Pi_{\mathcal{G}'\cap\mathcal{F}_t} = \Pi_{\mathcal{G}\vee\mathcal{G}'} \wedge \Pi_{\mathcal{F}_t} \textit{ for all } t \in \mathcal{T}.$$

Proof: For any partitions Π and Π', $\Pi \subset \Pi'$ if and only if $\Pi = \Pi'$. Thus, it suffices to show that

$$\Pi_{\mathcal{G}\cap\mathcal{F}_t} \vee \Pi_{\mathcal{G}'\cap\mathcal{F}_t} \subset \Pi_{\mathcal{G}\vee\mathcal{G}'} \wedge \Pi_{\mathcal{F}_t} \text{ for all } t \in \mathcal{T}.$$

Fix some $t \in \mathcal{T}$ and an event $A \in \Pi_{\mathcal{G}\cap\mathcal{F}_t} \vee \Pi_{\mathcal{G}'\cap\mathcal{F}_t}$. By definition of the supremum, $A = B \cap B'$ for some sets $B \in \Pi_{\mathcal{G}\cap\mathcal{F}_t}$ and $B' \in \Pi_{\mathcal{G}'\cap\mathcal{F}_t}$. Since $\mathcal{G}$ is connected, $\Pi_{\mathcal{G}\cap\mathcal{F}_t}\backslash\Pi_{\mathcal{F}_t} \subset \Pi_{\mathcal{G}\cap\mathcal{F}_T} = \Pi_{\mathcal{G}}$. Equivalently,

$$\Pi_{\mathcal{G}\cap\mathcal{F}_t} \subset \Pi_{\mathcal{G}} \cup \Pi_{\mathcal{F}_t}. \tag{6.12}$$

An analogous argument holds for $\mathcal{G}'$. By (6.12), there are two cases to consider:

If $B \in \Pi_{\mathcal{G}}$ and $B' \in \Pi_{\mathcal{G}'}$, then

$$B \cap B' \in (\Pi_{\mathcal{G}} \vee \Pi_{\mathcal{G}'}) \cap \mathcal{F}_t = \Pi_{\mathcal{G}\vee\mathcal{G}'} \cap \mathcal{F}_t \subset \Pi_{\mathcal{G}\vee\mathcal{G}'} \wedge \Pi_{\mathcal{F}_t}.$$

If $B \in \Pi_{\mathcal{F}_t}$ (or $B' \in \Pi_{\mathcal{F}_t}$), then $B \cap B' \in \Pi_{\mathcal{G}\cap\mathcal{F}_t} \vee \Pi_{\mathcal{G}'\cap\mathcal{F}_t} \leq \Pi_{\mathcal{F}_t}$ implies that $B = B \cap B'$. But then

$$B \cap B' = B \in \Pi_{\mathcal{F}_t} \cap \mathcal{G} \subset \Pi_{\mathcal{F}_t} \cap (\mathcal{G} \vee \mathcal{G}') \subset \Pi_{\mathcal{F}_t} \wedge \Pi_{\mathcal{G}\vee\mathcal{G}'}.\blacksquare$$

By Lemma 23, it is enough to prove that the class of sequentially connected algebras is a lattice. Take the supremum $\mathcal{G} \vee \mathcal{G}'$ of such algebras $\mathcal{G}$ and $\mathcal{G}'$ and an event $A \in \Pi_{(\mathcal{G}\vee\mathcal{G}')\cap\mathcal{F}_t}\backslash\Pi_{\mathcal{F}_t}$. By Lemma 23,

$$\begin{aligned}\Pi_{(\mathcal{G}\vee\mathcal{G}')\cap\mathcal{F}_t} &= \Pi_{\mathcal{G}\vee\mathcal{G}'} \wedge \Pi_{\mathcal{F}_t} \\ &= \Pi_{\mathcal{G}\cap\mathcal{F}_t} \vee \Pi_{\mathcal{G}'\cap\mathcal{F}_t}\end{aligned}$$

Thus there exist sets $B \in \Pi_{\mathcal{G}\cap\mathcal{F}_t}$ and $B' \in \Pi_{\mathcal{G}'\cap\mathcal{F}_t}$ such that $A = B \cap B'$. If $B \in \Pi_{\mathcal{F}_t}$, then $B' \in \mathcal{F}_t$ implies $A = B \cap B' = B \in \Pi_{\mathcal{F}_t}$ contradicting the choice of $A \notin \Pi_{\mathcal{F}_t}$. A symmetric argument implies $B' \in \Pi_{\mathcal{G}'\cap\mathcal{F}_t}\backslash\Pi_{\mathcal{F}_t}$. Since $\mathcal{G}$ and $\mathcal{G}'$ are sequentially connected, $B \in \Pi_{\mathcal{G}\cap\mathcal{F}_\tau}$ and $B' \in \Pi_{\mathcal{G}'\cap\mathcal{F}_\tau}$ for all $\tau \geq t$.

Conclude that $A = B \cap B' \in \Pi_{\mathcal{G} \cap \mathcal{F}_\tau} \vee \Pi_{\mathcal{G}' \cap \mathcal{F}_\tau} = \Pi_{\mathcal{G} \vee \mathcal{G}'} \wedge \Pi_{\mathcal{F}_\tau} = \Pi_{(\mathcal{G} \vee \mathcal{G}') \cap \mathcal{F}_\tau}$ for all $\tau \geq t$.

To show that $\mathcal{G} \wedge \mathcal{G}'$ is sequentially connected, take $A \in \Pi_{\mathcal{G} \cap \mathcal{G}' \cap \mathcal{F}_t} \backslash \Pi_{\mathcal{F}_t}$. Notice that

$$\begin{aligned} \Pi_{\mathcal{G} \cap \mathcal{G}' \cap \mathcal{F}_t} &= \Pi_{(\mathcal{G} \cap \mathcal{F})_t \wedge (\mathcal{G}' \cap \mathcal{F}_t)} = \\ &= \Pi_{\mathcal{G} \cap \mathcal{F}_t} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_t}. \end{aligned}$$

But $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_t} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_t}$ if and only if for all $B_0 \in \Pi_{\mathcal{G} \cap \mathcal{F}_t} \cup \Pi_{\mathcal{G}' \cap \mathcal{F}_t}$ such that $B_0 \subset A$:

$$A = \cup_{\{B_0, B_1, ..., B_k\}} \cup_{B \in \{B_0, B_1, ..., B_k\}} B,$$

where the union is taken over all sequences $\{B_0, B_1, ..., B_k\}$ of subsets of A such that consecutive elements intersect and belong alternatively to $\Pi_{\mathcal{G} \cap \mathcal{F}_t}$ and $\Pi_{\mathcal{G}' \cap \mathcal{F}_t}$.

Since $A \in (\Pi_{\mathcal{G} \cap \mathcal{F}_t} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_t}) \backslash \Pi_{\mathcal{F}_t}$, there exists a set $B_0 \in (\Pi_{\mathcal{G} \cap \mathcal{F}_t} \cup \Pi_{\mathcal{G}' \cap \mathcal{F}_t}) \backslash \Pi_{\mathcal{F}_t}$ such that $B_0 \subset A$. Fix such a set B_0 and consider a sequence $\{B_0, B_1, ..., B_k\}$ satisfying the conditions above. For $i < k$, each B_i intersects two disjoint subsets of $\mathcal{F}_t$ and so $B_i \notin \Pi_{\mathcal{F}_t}$. Moreover, if $B_k \in \Pi_{\mathcal{F}_t}$ then $B_k \cap B_{k-1} \neq \varnothing$ implies that $B_k \subset B_{k-1}$. Conclude that A can be written as the union over sequences $\{B_0, B_1, ..., B_k\}$ in $(\Pi_{\mathcal{G} \cap \mathcal{F}_t} \cup \Pi_{\mathcal{G}' \cap \mathcal{F}_t}) \backslash \Pi_{\mathcal{F}_t}$. Since $\mathcal{G}$ and $\mathcal{G}'$ are sequentially connected, A can be written as the union over sequences $\{B_0, B_1, ..., B_k\}$ in $\Pi_{\mathcal{G} \cap \mathcal{F}_\tau} \cup \Pi_{\mathcal{G}' \cap \mathcal{F}_\tau}$ for all $\tau \geq t$. Conclude that A must be a subset of some element in $\Pi_{\mathcal{G} \cap \mathcal{F}_\tau} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_\tau}$. But since $\Pi_{\mathcal{G} \cap \mathcal{F}_\tau} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_\tau}$ is finer than $\Pi_{\mathcal{G} \cap \mathcal{F}_t} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_t}$ and $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_t} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_t}$, it must be that $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_\tau} \wedge \Pi_{\mathcal{G}' \cap \mathcal{F}_\tau}$ for all $\tau \geq t$.■

Proposition 24 *A stopping time g induces a sequentially connected algebra.*

A stopping time is a function $g : \Omega \to \mathcal{T}$ such that $[g = t] \in \mathcal{F}_t$ for all $t \in \mathcal{T}$. The stopping time g is induces the algebra $\mathcal{G}$:

$$\mathcal{G} := \{A \in \mathcal{F}_T : A \cap [g = t] \in \mathcal{F}_t, \forall t \in \mathcal{T}\}.$$

To see that $\mathcal{G}$ is sequentially connected, first prove that

$$\{A \in \Pi_{\mathcal{F}_t} : A \subset [g = t]\} \subset \Pi_{\mathcal{G}}, \forall t \in \mathcal{T} \tag{6.13}$$

Fix $t \in \mathcal{T}$ and $A \in \Pi_{\mathcal{F}_t}$ such that $A \subset [g = t]$. Since $A \cap [g = t] = A \in \mathcal{F}_t$ and $A \cap [g = t'] = \varnothing \in \mathcal{F}_t$ for all $t' \neq t$, the definition of $\mathcal{G}$ implies $A \in \mathcal{G}$. If $B \subsetneq A \in \Pi_{\mathcal{F}_t}$, then $B \cap [g = t] = B \notin \mathcal{F}_t$ and thus $B \notin \mathcal{G}$. Conclude that $A \in \Pi_{\mathcal{G}}$.

Next fix $t < T$ and take $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_t} \backslash \Pi_{\mathcal{F}_t}$. If $A \cap [g \geq t] \neq \varnothing$, then there exists $A' \subsetneq A$ such that $A' \in \Pi_{\mathcal{F}_t}$ and $A' \cap [g \geq t] \neq \varnothing$. But $[g \geq t] = [g < t]^c \in \mathcal{F}_t$ and $A' \in \Pi_{\mathcal{F}_t}$ imply that $A' \subset [g \geq t]$ and so $A' \in \mathcal{G}$. In turn, $A' \in \mathcal{G} \cap \Pi_{\mathcal{F}_t}$ implies $A' \in \Pi_{\mathcal{G} \cap \mathcal{F}_t}$ contradicting the choice of A. Conclude that $A \subset [g < t]$.

Fix some $t' < t$ such that $A \cap [g = t'] \neq \varnothing$. Since $[g = t'] \in \mathcal{G} \cap \mathcal{F}_t$ and $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_t}$, it must be the case that $A \subset [g = t']$. This implies $A \in \mathcal{F}_{t'}$, for otherwise, $A \cap [g = t'] = A \notin \mathcal{F}_{t'}$ contradicts $A \in \mathcal{G}$. For any $A' \in \Pi_{\mathcal{F}_{t'}}$ and $A' \subset A \subset [g = t']$, equation (6.13) implies that $A' \in \Pi_{\mathcal{G}}$ and so $A' \in \Pi_{\mathcal{G} \cap \mathcal{F}_t}$. Since $A \in \Pi_{\mathcal{G} \cap \mathcal{F}_t}$, it must be the case that $A = A' \in \Pi_{\mathcal{G}}$. But then $A \in \Pi_{\mathcal{G}} \cap \mathcal{F}_{t+1} \subset \Pi_{\mathcal{G} \cap \mathcal{F}_{t+1}}$ as desired.■

The next example translates the Gabaix and Laibson [8] procedure for simplifying decision trees in the setting of this paper and shows that it induces a sequentially connected filtration.

Example 7 *(Satisficing) "Start from the initial node and follow all branches whose probability is greater than or equal to some threshold level α. Continue in this way down the tree. If a branch has a probability less than α, consider the node it leads to, but do not advance beyond that node."*

Thus, let μ be a measure on $(\Omega, \mathcal{F}_T)$ and $\alpha \in [0,1]$ be a threshold level. For each event A and algebra $\mathcal{F}$, define $r_{\mathcal{F}}(A)$ to be the smallest $\mathcal{F}$-measurable superset of A. The collection of events $\{\mathcal{A}_t\}$ is a satisficing procedure if:

$$
\begin{aligned}
\mathcal{A}_t &= \{\Pi_{\mathcal{F}_t}\} \text{ for all } t \leq 1, \text{ and for all } t > 1 \\
\mathcal{A}_t &= \{A \in \Pi_{\mathcal{F}_t} : r_{\mathcal{F}_{t-1}}(A) \in \mathcal{G}_{t-1} \text{ and } \mu(r_{\mathcal{F}_{t-1}}(A) \mid r_{\mathcal{F}_{t-2}}(A)) \geq \alpha\}.
\end{aligned}
$$

It is not difficult to see that $\{\mathcal{A}_t\}$ generates a sequentially connected filtration. In fact, the filtration is induced by the stopping time:

$$
\begin{aligned}
[\tau &\leq 1] = \varnothing, \text{ and for all } t > 1 \\
[\tau &= t] = \cup \left\{ A \in \Pi_{\mathcal{F}_t} : \mu(A \mid r_{\mathcal{F}_{t-1}}(A)) < \alpha \text{ and } \mu(r_{\mathcal{F}_{t-1}}(A) \mid r_{\mathcal{F}_{t-2}}(A)) \geq \alpha \right\}.
\end{aligned}
$$

References

[1] C. Aliprantis and K. Border. *Infinite Dimensional Analysis*. Springer-Verlag, 2 edition, 1999.

[2] F. Anscombe and R. Aumann. A definition of subjective probability. *The Annals of Mathematical Statistics*, 34:199–205, 1963.

[3] E. Dekel, B. Lipman, and A. Rustichini. Recent developments in modeling unforeseen contingencies. *European Economic Review*, 42:523–542, 1998.

[4] P. Diaconis and S. Zabell. Updating subjective probability. *Journal of the American Statistical Association*, 77:822–830, 1982.

[5] L. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113:1–31, 2003.

[6] L. Epstein and S. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

[7] L. Epstein, M. Marinacci, and K. Seo. Coarse contingencies and ambiguity. *Theoretical Economics*, 2(4):355–394, 2007.

[8] X. Gabaix and D. Laibson. A boundedly rational decision algorithm. *The American Economic Review*, 90(2):433–438, May 2000.

[9] P. Ghirardato. Coping with ignorance: Unforeseen contingencies and non-additive uncertainty. *Economic Theory*, 17(2):247–276, 2001.

[10] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

[11] I. Gilboa and D. Schmeidler. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research*, 52: 43–65, 1994.

[12] P. Klibanoff, M. Marinacci, and S. Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144:930–976, 2009.

[13] D. Kreps. Static choice in the presence of unforeseen contingencies. In P. Dasgupta, D. Gale, O. Hart, and E. Maskin, editors, *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, pages 259–281. MIT Press, Cambridge, MA, 1992.

[14] D. Kreps and E. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978.

[15] S. Mukerji. Understanding the nonadditive probability decision model. *Economic Theory*, 9:23–46, 1997.

[16] L. Savage. *The Foundations of Statistics*. Wiley, New York, NY, 1954.

[17] O. Williamson. *The Economic Institutions of Capitalism*. The Free Press, 1985.

[18] S. Zabell. Predicting the unpredictable. *Synthese*, 90:205–232, 1992.