

Non-Bayesian updating: a theoretical framework

LARRY G. EPSTEIN

Department of Economics, Boston University

JAWWAD NOOR

Department of Economics, Boston University

ALVARO SANDRONI

Department of Economics, University of Pennsylvania

This paper models an agent in a multi-period setting who does not update according to Bayes' Rule, and who is self-aware and anticipates her updating behavior when formulating plans. Choice-theoretic axiomatic foundations are provided to capture updating biases that reflect excessive weight given to either prior beliefs, or, alternatively, to observed data. A counterpart of the exchangeable Bayesian learning model is also described.

KEYWORDS. Non-Bayesian updating, temptation and self-control, overreaction, underreaction, learning, law of small numbers.

JEL CLASSIFICATION. D81.

1. INTRODUCTION

Epstein (2006) models an agent who does not update according to Bayes' Rule, but is self-aware and anticipates her updating behavior when formulating plans. He provides axiomatic foundations for his model in the form of a representation theorem for suitably defined preferences such that *both* the prior *and* the way in which it is updated are subjective. The model is nested in a three-period framework, where the agent updates once and consumption occurs only at the terminal time. This paper extends the model to an infinite horizon setting, thereby enabling it to address dynamic issues and making it more amenable to applications.

The benchmark for the present model is the standard specification of utility in dynamic modeling, whereby utility at time t is given by

$$U_t(c) = E_t \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} u(c_\tau) \right] \quad t = 0, 1, \dots, \quad (1)$$

Larry G. Epstein: lepstein@bu.edu

Jawwad Noor: jnoor@bu.edu

Alvaro Sandroni: sandroni@econ.upenn.edu

We have benefitted from comments from two referees and an editor, and from conversations with Mark Machina, Massimo Marinacci, Joseph Perktold, Werner Ploberger, and especially Igor Kopylov.

where $c = (c_\tau)$ is a consumption process, δ and u have the familiar interpretations, and E_t denotes the expectation operator associated with a subjective prior updated by Bayes' Rule. Our model generalizes (1), to which it reduces when updating conforms to Bayes' Rule.

The model shares some similarities with the Gul and Pesendorfer (2001, 2004) model of temptation and self-control.¹ While these authors (henceforth GP) focus on behavior associated with non-geometric discounting, we adapt their approach to model non-Bayesian updating. The connection drawn here between temptation and updating is as follows: at period t , the agent has a prior view of the relationship between the next observation s_{t+1} and the future uncertainty $(s_{t+2}, s_{t+3}, \dots)$ that she considers 'correct.' But after observing a particular realization s_{t+1} , she changes her view on the noted relationship. For example, she may respond exuberantly to a good (or bad) signal after it is realized and decide that it is an even better (or worse) signal about future states than she had thought ex ante. She tries to resist the temptation to behave in accordance with the new view rather than in accordance with the view she considers correct. Temptation might be resisted but at a cost. Thus she acts as though forming a compromise posterior belief; it differs from what would be implied by Bayesian updating of the original prior and in that sense reflects non-Bayesian updating. The exuberant agent described above would appear to an outside observer as someone who overreacts to data.

An important feature of our model from the point of view of applications is its richness: just as the Savage and Anscombe–Aumann theorems provide foundations for subjective expected utility theory without restricting beliefs, the present framework imposes a specific structure for preferences without unduly restricting the nature of updating. We demonstrate richness by describing specializations that capture excessive weight given, at the updating stage, to prior beliefs, or alternatively, to the observed data. In addition, we describe a counterpart of the exchangeable Bayesian learning model.

To illustrate the scope of our framework, consider an agent who is trying to learn the true parameter in a set Θ . Updating of beliefs in response to observations s_1, \dots, s_t , leads to the process of posteriors $\{\mu_t\}$, where each μ_t is a probability measure on Θ . Bayesian updating leads to the process

$$\mu_{t+1} = BU(\mu_t; s_{t+1}),$$

where $BU(\mu_t; s_{t+1})$ denotes the Bayesian update of μ_t . One alternative consistent with our model is the process

$$\mu_{t+1} = (1 - \kappa_{t+1})BU(\mu_t; s_{t+1}) + \kappa_{t+1}\mu_t,$$

where $\kappa_{t+1} \leq 1$. If κ_{t+1} does not depend on the latest observation s_{t+1} and if $\kappa_{t+1} > 0$, then the updating rule can be interpreted as attaching too much weight to prior beliefs μ_t and hence underreacting to observations. Another alternative has the form

$$\mu_{t+1} = (1 - \kappa_{t+1})BU(\mu_t; s_{t+1}) + \kappa_{t+1}\psi_{t+1},$$

¹At a technical level, we rely heavily on generalizations of the Gul–Pesendorfer model proved by Kopylov (2007).

where ψ_0 is a suitable noninformative prior and subsequent ψ_t 's are obtained via Bayesian updating. This updating rule for the posteriors μ_t can be interpreted (under the assumptions for κ_{t+1} stated above) as attaching too much weight to the sample.²

Several systematic deviations from Bayesian updating have been observed in experimental psychology; see [Tversky and Kahneman \(1974\)](#) and the surveys by [Camerer \(1995\)](#) and [Rabin \(1998\)](#), for example. This evidence deals with the updating of objective probabilities. Thus models such as [Rabin \(2002\)](#) and [Mullainathan \(2000\)](#), for example, that address the experimental evidence take probabilities as directly observable. In contrast, we follow the Savage tradition and address the seemingly more relevant case where probabilities are subjective; indeed, our model of (or story about) updating, is more intuitive if probabilities are subjective. This forces us to focus on behavior, in the form of axioms on preferences, that reveals both beliefs and updating. (The cited models that assume objective probabilities are not explicit about the associated model of choice.) Though our model does not address the experimental evidence directly, the two nevertheless are related. This is because one suspects that some of the biases noted in the experimental literature would be exhibited also when updating subjective probabilities, and because, as will become evident, our framework is rich enough to accommodate a wide range of deviations from Bayesian updating.

The paper proceeds as follows. [Section 2](#) defines the formal domain of choice, the space of contingent menus, and then functional forms for conditional utility functions. Some specializations corresponding to specific updating biases and to learning about parameters are described in [Section 3](#). Finally, axiomatic foundations are provided in [Section 4](#). Proofs are collected in the [Appendix](#).

2. UTILITY

2.1 Primitives

Time is discrete and varies over $t = 0, 1, 2, \dots$. Uncertainty is represented by a (finite) period state space S , one element of which is realized at each t . Thus the complete uncertainty is represented by the full state space $\times_{t=1}^{\infty} S_t$, where $S_t = S$ for all $t > 0$. The period consumption space is $C_t = C$, a compact metric mixture space.³ Though we often refer to c_t in C_t as period t consumption, it is more accurately thought of as a lottery over period t consumption. Thus we adopt an Anscombe–Aumann style domain where outcomes are lotteries. Information available at t is given by the history $s_1^t = (s_1, \dots, s_t)$. Thus time t consumption, conditional beliefs, conditional preferences, and so on, are taken to be suitably measurable, though dependence on s_1^t is often suppressed in the notation.

²While this paper focuses on presenting the framework, in [Epstein et al. \(2008\)](#) we apply the framework to ask: *what do non-Bayesian updaters learn?* We show, for instance, that multiple repetitions of non-Bayesian updating rules that underreact to observations uncover the true data generating process with probability one, while non-Bayesian updaters who overreact can, with positive probability, become certain that a false parameter is true and thus converge to incorrect forecasts.

³We use this term to include the property that the mixture operation $(c, c', \alpha) \mapsto \alpha c + (1 - \alpha)c'$ is continuous with respect to the product metric on $C \times C \times [0, 1]$.




FIGURE 1. The time line.

For any compact metric space X , the set of acts from S into X is X^S ; it is endowed with the product topology. A closed (hence compact) subset of $C \times X^S$ is called a *menu* (of pairs (c, F) , where $c \in C$ and $F \in X^S$). Denote by $\mathcal{M}(X)$ the set of all compact subsets of X , endowed with the Hausdorff metric. Analogously, $\mathcal{M}(C \times X^S)$ is the set of menus of pairs (c, F) as above; it inherits the compact metric property (Aliprantis and Border 1994, Section 3.16).

Consider a physical action taken at time t , where consumption at t has already been determined. The consequence of that action is a menu, contingent on the state s_{t+1} , of alternatives for $t + 1$, where these alternatives include both choices to be made at $t + 1$ —namely, the choice of both consumption and also another action. This motivates identifying each physical action with a *contingent menu*, denoted F , where

$$F : S \longrightarrow \mathcal{M}(C \times \mathcal{C}), \tag{2}$$

and \mathcal{C} denotes the space of all contingent menus. The preceding suggests that \mathcal{C} can be identified with $(\mathcal{M}(C \times \mathcal{C}))^S$. Appendix A shows the existence of a (compact metric) \mathcal{C} satisfying the homeomorphism

$$\underset{\text{homeo}}{\mathcal{C}} \approx (\mathcal{M}(C \times \mathcal{C}))^S. \tag{3}$$

Hence, we identify any element of \mathcal{C} with a mapping F as in (2).

Though the domain \mathcal{C} is time stationary and applies at every t , when we wish to emphasize that a particular choice is made at t , we write that the agent chooses contingent menu $F_t \in \mathcal{C}_t$,

$$F_t : S_{t+1} \longrightarrow \mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1}), \tag{4}$$

where $\mathcal{C}_t = \mathcal{C}_{t+1} = \mathcal{C}$. (Keep in mind that we have previously defined $S_{t+1} = S$ and $C_{t+1} = C$.)

The final primitive is a process of preference relations $(\succeq_t)_{t=0}^\infty$, one for each time t and history s_t^1 , where the domain of \succeq_t is $C_t \times \mathcal{C}_t$. At time 0, the agent uses \succeq_0 to choose (c_0, F_0) in $C_0 \times \mathcal{C}_0$. She does this as though anticipating the following: at 1^- , a signal s_1 is realized, and this determines a menu $F_0(s_1) \subset C_1 \times \mathcal{C}_1$; at time 1, she updates and uses the order \succeq_1 (which corresponds to the history s_1) to choose some (c_1, F_1) from $F_0(s_1)$. She consumes c_1 and her (contingent) options for the future are described by F_1 . Continuing in this way, and given some previous choice of contingent menu F_t , she observes a signal s_{t+1} , updates and then uses the order \succeq_{t+1} (corresponding to the history $(s_1, s_2, \dots, s_{t+1})$) to choose some (c_{t+1}, F_{t+1}) from $F_t(s_{t+1})$. (See Figure 1.)

This completes the description of the primitives and the setting. Before presenting the formal model we outline the story behind it with the help of the above time line.

At time 0 the agent formulates a prior over the full state space $\times_{t=1}^{\infty} S_t$. She uses this to evaluate any alternative (c_0, F_0) , which describes contingent options. She is forward-looking and so her evaluation of (c_0, F_0) also takes into consideration what choice she expects to make from $F_0(s_1)$ at time 1, for each s_1 . She anticipates that her choice will be subject to temptation: she will be tempted at that time to deviate from the view of the world she possesses at time 0, and to use a ‘temptation belief’ to guide her choice. Depending on how successful she is at exerting self-control, she will end up using a compromise belief to guide her choice at time 1: this is a mixture of the temptation belief and the Bayesian update of her time 0 belief. To an outside observer, she is thus not a Bayesian updater. At time 1 she will be in the same position she was at time 0, possessing some view of the world and anticipating a struggle at time 2 with temptation to deviate from this view. So on and so forth for all t .

2.2 Functional form

We describe the representation of (\succeq_t) ; axiomatic foundations are deferred to [Section 4](#). Components of the functional form include: a discount factor $0 < \delta < 1$, $u : C \rightarrow \mathbb{R}^1$ linear, continuous, and nonconstant, a probability measure p_0 on S_1 with full support, and an adapted process $(p_t, q_t, \alpha_t)_{t=1}^{\infty}$, where,⁴

$$\alpha_t \in (0, 1], p_t, q_t \in \Delta(S_{t+1}), \text{ and each } p_t \text{ has full support.}$$

For each $(c_t, F_t) \in C_t \times \mathcal{C}_t$, define

$$\mathcal{U}_t(c_t, F_t) = u(c_t) + \delta \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dp_t, \quad t \geq 0, \tag{5}$$

$$\mathcal{V}_t(c_t, F_t) = u(c_t) + \delta \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dq_t, \quad t > 0, \tag{6}$$

where $U_{t+1}(\cdot, s_{t+1}) : \mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1}) \rightarrow \mathbb{R}^1$ is defined recursively by⁵

$$U_{t+1}(M, s_{t+1}) = \max_{(c_{t+1}, F_{t+1}) \in M} \left\{ \mathcal{U}_{t+1}(c_{t+1}, F_{t+1}) + \frac{1 - \alpha_{t+1}}{\alpha_{t+1}} (\mathcal{V}_{t+1}(c_{t+1}, F_{t+1}) - \max_{(c'_{t+1}, F'_{t+1}) \in M} \mathcal{V}_{t+1}(c'_{t+1}, F'_{t+1})) \right\}. \tag{7}$$

Then \succeq_0 is represented by $\mathcal{U}_0(\cdot)$ and for each $t > 0$, \succeq_t is represented by $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$.

The Bayesian intertemporal utility model (1) is specified by u , δ , and a process (p_t) of one-step-ahead conditionals, which determines a unique prior on the full state space

⁴ $\Delta(S)$ is the set of probability measures on the finite set S . A stochastic process (X_t) on $\times_1^{\infty} S_t$ is adapted if X_t is measurable with respect to the σ -algebra \mathcal{S}_t that is generated by all sets of the form $\{s_1\} \times \dots \times \{s_t\} \times \times_{t+1}^{\infty} S_t$. Below we often write $p_t(\cdot)$ rather than $p_t(\cdot | s_t^t)$. When we want to emphasize dependence on the last observation s_t , we write $p_t(\cdot | s_t)$. Similarly, history is suppressed from our notation when we write $\mathcal{U}_t(c_t, F_t)$ and $\mathcal{V}_t(c_t, F_t)$.

⁵See [Theorem 1](#) for conditions under which utility is well-defined by this recursion.

$\times_1^\infty S_t$. It is obtained as the special case where $(1 - \alpha_t)(q_t - p_t) \equiv 0$ for all t . Then (7) reduces to

$$U_{t+1}(M, s_{t+1}) = \max_{(c_{t+1}, F_{t+1}) \in M} \mathcal{U}_{t+1}(c_{t+1}, F_{t+1}),$$

and \succeq_t is represented by

$$\mathcal{U}_t(c_t, F_t) = u(c) + \delta \int_{S_{t+1}} \left(\max_{(c_{t+1}, F_{t+1}) \in F_{t+1}(s_{t+1})} \mathcal{U}_{t+1}(c_{t+1}, F_{t+1}) \right) dp_t, \quad (c_t, F_t) \in C_t \times \mathcal{C}_t.$$

This is the standard model in the sense that it extends the model of utility over consumption processes given by (1) to contingent menus by assuming that menus are valued according to the best alternative they contain (a property termed *strategic rationality* by Kreps 1988). In particular, time t conditional beliefs about the future are obtained by applying Bayes' Rule to the prior on $\times_1^\infty S_t$ that is induced by the one-step-ahead conditionals (p_t).

More generally, *two* processes of one-step-ahead conditionals, p_t 's and q_t 's, must be specified, as well as the process of α_t 's. The way in which these deliver non-Bayesian updating is explained below along with further discussion and interpretation. Sections 3.1 and 3.2 provide several examples. See also Epstein (2006) for discussion in the context of a three-period model.

2.3 Interpretation

To facilitate interpretation, and also for later purposes, consider some subclasses of \mathcal{C}_t . The contingent menu F_t provides commitment for the next period if $F_t(s_{t+1})$ is a singleton for each s_{t+1} . The set of contingent menus that provide commitment for *all* future periods is denoted by $\mathcal{C}_t^c = \mathcal{C}^c \subset \mathcal{C}$. Each F_t in \mathcal{C}_t^c determines a unique (random variable) consumption process $c^{F_t} = (c_\tau^{F_t})_{\tau \geq t}$. If each $c_\tau^{F_t}$ is measurable with respect to information at time $t + 1$, then all uncertainty is resolved next period; the set of all such contingent menus is $\mathcal{C}_t^{c,+1} = \mathcal{C}^{c,+1} \subset \mathcal{C}^c$.⁶ An example is a (one-step-ahead) *bet on the event* $G \subset S_{t+1}$, which pays off with a good deterministic consumption stream if the state next period lies in G and with a poor one otherwise.

For any c_t and contingent menu F_t that provides commitment ($F_t \in \mathcal{C}_t^c$), we have

$$\mathcal{U}_t(c, F) = u(c_t) + \delta \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dp_t(s_{t+1}).$$

It follows that if $F \in \mathcal{C}_0^c$, then

$$\mathcal{U}_0(c, F) = \int_{S_1 \times S_2 \times \dots} \sum_1^\infty \delta^{t-1} u(c_t^F) dP_0(\cdot),$$

where c^F is the consumption process induced by F as just explained, and $P_0(\cdot)$ is the unique measure on $\times_1^\infty S_t$ satisfying, for every T ,

$$P_0(s_1, s_2, \dots, s_{T+1}) = p_0(s_1) \times \dots \times p_t(s_{t+1} | s_1^t) \times \dots \times p_T(s_{T+1} | s_1^T).$$

⁶See Appendix A for some formal details regarding \mathcal{C}^c and $\mathcal{C}^{c,+1}$.

Thus \succeq_0 restricted to \mathcal{C}_0^c conforms to subjective expected (intertemporally additive) utility with prior P_0 . The ranking of commitment prospects at 0 leaves no choices to be made later and thus reveals nothing about future updating— P_0 reflects only an ex ante view.

To interpret P_0 further, consider its one-step-ahead conditionals p_t for $t \geq 1$. Because these conditional beliefs are formed for contingencies that are ‘distant’ (at least two periods ahead), they are based on a degree of detachment and objectivity and thus the agent views them as ‘correct.’⁷ She continues to view them as correct as time passes. If she were not subject to other influences, her posterior at t would be

$$P_t(s_{t+1}, s_{t+2}, \dots, s_{T+1} | s_1^t) = p_t(s_{t+1} | s_1^t) \times \dots \times p_T(s_{T+1} | s_1^T),$$

the Bayesian update of P_0 . However, as explained shortly, she may update differently and be led to different posteriors.

Her actual updating underlies the preference \succeq_t prevailing after an arbitrary history s_1^t . By assumption, \succeq_t is represented by $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$. To proceed, define the one-step-ahead conditional measure m_t by: $m_0 = p_0$ and, for $t > 0$,

$$m_t(s_{t+1}) = m_t(s_{t+1} | s_1^t) = \alpha_t p_t(s_{t+1} | s_1^t) + (1 - \alpha_t) q_t(s_{t+1} | s_1^t).$$

Next, compute that for any c_t and any contingent menu $F_t \in \mathcal{C}_t^c$ that provides commitment for periods beyond t ,

$$\begin{aligned} \alpha_t \mathcal{U}_t(c_t, F_t) + (1 - \alpha_t) \mathcal{V}_t(c_t, F_t) &= u(c_t) + \delta \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dm_t(s_{t+1}) \\ &= \int_{S_{t+1} \times S_{t+2} \times \dots} \sum_{t+1}^{\infty} \delta^{\tau-t-1} u(c_\tau^{F_t}) dQ_t(\cdot | s_1^t), \end{aligned}$$

where $Q_t(\cdot | s_1^t)$ is the unique measure on $\times_{t+1}^{\infty} S_\tau$ satisfying, for every T ,

$$Q_t(s_{t+1}, s_{t+2}, \dots, s_{T+1} | s_1^t) = m_t(s_{t+1} | s_1^t) \times p_{t+1}(s_{t+2} | s_1^{t+1}) \times \dots \times p_T(s_{T+1} | s_1^T).$$

Evidently, at t the agent’s behavior (at least within \mathcal{C}_t^c) corresponds to the posterior $Q_t(\cdot | s_1^t)$, and this differs from the period 0 perspective $P_t(\cdot | s_1^t)$. Note that Q_t is not the Bayesian update of P_0 , nor is it the Bayesian update of Q_{t-1} . The difference between P_t and Q_t lies in the way that one-step-ahead beliefs over S_{t+1} are formulated—the conditional one-step-ahead belief actually adopted at t is $m_t(\cdot)$, whereas the one that seems appropriate from the perspective of the initial period is $p_t(\cdot)$.⁸

⁷Since p_0 is not relevant to the subsequent response to signals, its interpretation is less important here. See the comments at the end of the section.

⁸The behavioral meaning of m_t is sharper if we restrict attention to contingent menus in $\mathcal{C}_t^{c,+1}$ (providing perfect commitment and such that all uncertainty resolves at $t + 1$). Then beliefs about states in $S_{t+2} \times S_{t+3} \times \dots$ are irrelevant, so we conclude that m_t guides the ranking of such contingent menus; for example, it guides the ranking of bets on S_{t+1} . Because the ranking of one-step-ahead bets, and more specifically the way in which it depends on past observations, is a common and natural way to understand updating behavior, we refer to m_t frequently below when considering more specific models.

The story underlying the noted difference between P_t and Q_t is as follows: consider the evaluation of a pair (c_t, F_t) in $C_t \times \mathcal{C}_t$ after having observed the history s_t^t . The functions \mathcal{U}_t and \mathcal{V}_t describe two ways that (c_t, F_t) may be evaluated. Both evaluate immediate consumption c_t in the same way, and they discount the expected utility of the contingent menu F_t in the same way as well. However, they disagree on how to compute the expected utility of F_t : \mathcal{U}_t uses p_t and \mathcal{V}_t uses q_t . The former is the ‘correct’ one-step-ahead conditional. But in our model, after having observed s_t^t , the agent changes her view of the world to the one-step-ahead conditional q_t . For instance, if s_1^t represents a run of bad signals, she may believe that the likelihood of another bad state is higher than her ex ante assessment. Alternatively, she may feel that a good signal ‘is due’ and thus assign it a higher conditional probability than she did when anticipating possibilities with the cool-headedness afforded by temporal distance. Thus there are conflicting incentives impinging on the agent at t . The period 0 perspective calls for maximizing \mathcal{U}_t , but having seen the sample history s_t^t and having changed her view of the world, she is tempted to maximize \mathcal{V}_t . Resisting temptation is costly and she recognizes that the time 0 perspective is ‘correct.’ She is led to compromise and to maximize $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$, the utility function representing \succeq_t . The corresponding behavior is as though she used the *compromise* one-step-ahead conditional $\alpha_t p_t + (1 - \alpha_t) q_t$, which is just m_t . The parameter α_t captures her ability to resist temptation.

The cost of self-control incurred when compromising between $\mathcal{U}_t(\cdot)$ and $\mathcal{V}_t(\cdot)$ is reflected not in the representation of \succeq_t , but rather in that of \succeq_{t-1} , specifically in the utility of a menu $M_t \in \mathcal{M}(C_t \times \mathcal{C}_t)$ given by the function $U_t(M_t, s_t)$. The nonpositive term

$$\frac{1 - \alpha_t}{\alpha_t} \left[\mathcal{V}_t(c_t, F_t) - \max_{(c'_t, F'_t) \in M_t} \mathcal{V}_t(c'_t, F'_t) \right] \leq 0,$$

appearing in (7) can be interpreted as the utility cost of self-control. Thus (7) states that for any menu M_t received after the history s_t^t , $U_t(M_t, s_t)$ is the maximum over M_t of $\mathcal{U}_t(\cdot)$ net of self-control costs. Observe that this maximization is equivalent to

$$\max_{(c_t, F_t) \in M_t} \left\{ \mathcal{U}_t(\cdot) + \frac{1 - \alpha_t}{\alpha_t} \mathcal{V}_t(\cdot) \right\},$$

and that $\mathcal{U}_t(\cdot) + ((1 - \alpha_t)/\alpha_t) \mathcal{V}_t(\cdot)$ represents \succeq_t . Thus (7) suggests that choosing the \succeq_t -best element in M_t involves incurring a utility cost of self-control.

Unlike a standard agent, our agent may later deviate from her current view of conditional likelihoods. Indeed, she is “dynamically inconsistent” in the sense that she may not follow through with a committed contingent consumption plan in \mathcal{C}_t^c if somehow she has the opportunity to undo previous commitments. Our agent is also self-aware and forward looking—she anticipates at any time t that she will later adopt conditional beliefs different from those that seem correct now. Thus she may value commitment: a smaller menu may be strictly preferable because it could reduce self-control costs.⁹ In spite of the value of commitment, the above constitutes a coherent model of dynamic

⁹Indeed, since temptation arises only because of non-Bayesian updating, the agent exhibits a preference for commitment *if and only if* she is a non-Bayesian updaters.

choice. Unlike the case in the modeling approach growing out of [Strotz \(1955\)](#), there is no need to add assumptions about how the agent resolves her intertemporal inconsistencies. If you like, these resolutions are already embedded in her utility function defined on contingent menus. This aspect of the model uses the insight of GP.

A difference from GP is in terms of the primitives of the model. The primitive adopted by GP, and also by [Epstein \(2006\)](#), is a single preference ordering that describes choices at one point in time. A story about choices in subsequent periods is only “suggested” by the primitive preference and, in particular, its representation. In our model, the primitive consists of in-principle-observable preferences in each period.¹⁰ Foundations for our model thus specify the testable implications for dynamic choice, as opposed to implications only for period 0 preference as in GP and [Epstein \(2006\)](#).

Finally, a comment on the seeming asymmetry in the representations of \succeq_0 and \succeq_t for $t > 0$ is in order. The utility function $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$ for $t > 0$ makes explicit the conflict experienced by the agent in forming the belief m_t over S_{t+1} . The representation $\mathcal{U}_0(\cdot)$ for \succeq_0 is agnostic in this regard: it says nothing beyond the fact that at 0 the agent has some belief p_0 over S_1 , which may or may not have been formed after resolving some conflict. Thus the representations tell the same story, except that the decomposition of p_0 into its ‘correct’ and temptation components is not specified. The reason for the latter stems from the fact that, as in GP, we take a preference for commitment as the behavioral manifestation of a conflict—the decomposition of the belief m_t into its correct ‘ p_t ’ and temptation ‘ q_t ’ components is based on preferences, in particular on attitudes towards commitment opportunities, prevailing at time $t - 1$. A similar decomposition of p_0 would involve preferences in (unmodeled) periods prior to time 0. The reader should note, however, that p_0 is not relevant for understanding updating behavior, and consequently, its decomposition is of little interest for our purposes.

3. SOME SPECIFICATIONS

The framework described above is rich. One way to see this is to focus on one-step-ahead beliefs at any time $t + 1$. As pointed out in the previous section, these are represented by $m_{t+1} = \alpha_{t+1} p_{t+1} + (1 - \alpha_{t+1}) q_{t+1}$, while Bayesian updating of time t beliefs would lead to beliefs described by p_{t+1} . Thus, speaking roughly, updating deviates from Bayes’ Rule in a direction given by $q_{t+1} - p_{t+1}$ and to a degree determined by α_{t+1} , neither of which is constrained by our framework. Consequently, the modeler is free to specify the nature and degree of the updating bias, including how these vary with history, in much the same way that a modeler who works within the Savage or Anscombe–Aumann framework of subjective expected utility theory is free to specify beliefs as she sees fit. To illustrate, we describe specializations of the model that impose structure on updating. Two alternatives are explored, whereby excess weight at the updating stage is given to either (i) prior beliefs, or (ii) the sample frequency. In both cases, restrictions are imposed on the relation between q_{t+1} and p_{t+1} , but not on α_{t+1} ; thus they limit the direction but not the magnitude of the updating bias.¹¹ We also consider a specification

¹⁰In this respect, our model is related to [Noor \(2008\)](#).

¹¹Axiomatic characterizations of these specializations are given in [Section 4.2](#).




FIGURE 2. The time line.

of our model so as to capture the case where the data generating process is unknown up to a parameter.

3.1 Updating biases

The first specialization, termed *prior-bias*, corresponds to the restriction

$$q_{t+1}(\cdot | s_{t+1}) = (1 - \lambda_{t+1})p_{t+1}(\cdot | s_{t+1}) + \lambda_{t+1} \sum_{s'_{t+1}} m_t(s'_{t+1})p_{t+1}(\cdot | s'_{t+1}), \tag{8}$$

for some adapted process (λ_t) with $\lambda_{t+1} \leq 1$.¹² We use the term (i) *positive prior-bias* or (ii) *negative prior-bias* if (8) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$. Note that (8) defines all q_t 's inductively given the p_t 's and λ_t 's. Thus the corresponding model of utility is completely specified by δ, u, p_0 , and the process $(p_t, \alpha_t, \lambda_t)_{t \geq 1}$.

To interpret (8), think of the agent at time $t > 0$, after the history s_1^t has been realized, holding a view about $\times_{t+1}^\infty S_\tau$, and in particular about S_{t+2} . On observing the further realization s_{t+1} at $t + 1$, she forms new beliefs about S_{t+2} by updating this view. (See Figure 2.) The restriction (8) implies that when updating, she attaches inordinate weight to prior (time t) beliefs over S_{t+2} .

To see why, recall from the previous section that at t , after the history s_1^t , the agent's beliefs about future uncertainty are captured by the measure

$$Q_t(s_{t+1}, s_{t+2}, \dots, s_{T+1} | s_1^t) = m_t(s_{t+1})p_{t+1}(s_{t+2} | s_1^{t+1}) \times \dots \times p_T(s_{T+1} | s_1^T).$$

Refer to this measure as the agent's prior view at t . The measure

$$\sum_{s'_{t+1}} m_t(s'_{t+1})p_{t+1}(\cdot | s'_{t+1})$$

represents beliefs about S_{t+2} held at t ; refer to it as the prior view of S_{t+2} at t , while the measure $p_{t+1}(\cdot | s_{t+1})$ over S_{t+2} is the Bayesian update of the prior view at t conditional on observing s_{t+1} . If $\lambda_{t+1} = 0$ or $q_{t+1} = p_{t+1}$, then updating consists of responding to data by applying Bayes' Rule to the prior view. On the other hand, if $\lambda_{t+1} = 1$, then the prior view of S_{t+2} (expressed by $\sum_{s'_{t+1}} m_t(s'_{t+1})p_{t+1}(\cdot | s'_{t+1})$) is also the posterior, which gives all the weight to prior beliefs and none to data. Thus an agent who updates according to the average scheme in (8) exhibits a positive bias to the prior if $\lambda_{t+1} > 0$ and a negative bias if $\lambda_{t+1} < 0$.

¹²When $\lambda_{t+1} < 0$ in (10), q_{t+1} is well-defined as a probability measure only under special conditions.

Though q_{t+1} leads to urges for making choices at $t + 1$, the agent balances it with the view represented by p_{t+1} as described in Section 2.3, and acts as though she forms the compromise one-step-ahead posterior $m_{t+1} = \alpha_{t+1}p_{t+1} + (1 - \alpha_{t+1})q_{t+1}$. The above noted bias of q_{t+1} extends to this mixture of p_{t+1} and q_{t+1} : substitute for q_{t+1} from (8) and deduce that

$$m_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1}))p_{t+1} + \lambda_{t+1}(1 - \alpha_{t+1}) \sum_{s'_{t+1}} m_t(s'_{t+1})p_{t+1}(\cdot | s'_{t+1}), \quad (9)$$

which admits an interpretation analogous to that described above.¹³

Further content can be introduced into the model described in (8) by imposing structure on the way in which λ_{t+1} depends on the history s_1^{t+1} . For example, it might depend not only on the empirical frequency of observations but also on their order due to sensitivity to streaks or other patterns. While each specialization we have described fixes a sign for λ_{t+1} that is constant across times and histories, one can imagine that an agent might react differently depending on the history. Formulating a theory of the λ_{t+1} 's is a subject for future research.

Denote by Ψ_{t+1} the empirical frequency measure on S given the history s_1^{t+1} ; that is, $\Psi_{t+1}(s)$ is the relative frequency of s in the sample s_1^{t+1} . The second bias, termed *sample-bias*, corresponds to the restriction

$$q_{t+1}(\cdot | s_{t+1}) = (1 - \lambda_{t+1})p_{t+1}(\cdot | s_{t+1}) + \lambda_{t+1}\Psi_{t+1}(\cdot), \quad (10)$$

for some adapted process (λ_t) with $\lambda_{t+1} \leq 1$.¹⁴ We use the term (i) *positive sample-bias* or (ii) *negative sample-bias* if (10) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$.

The interpretation is similar to that for prior-bias. The implied adjustment rule for one-step-ahead beliefs is

$$m_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1}))p_{t+1} + \lambda_{t+1}(1 - \alpha_{t+1})\Psi_{t+1}.$$

Under positive sample-bias ($\lambda_{t+1} \geq 0$), the Bayesian update $p_{t+1}(s_{t+2})$ is adjusted in the direction of the sample frequency $\Psi_{t+1}(s_{t+2})$, implying a bias akin to the *hot-hand fallacy*—the tendency to over-predict the continuation of recent observations. For negative sample-bias,

$$m_{t+1} = p_{t+1} + (-\lambda_{t+1}(1 - \alpha_{t+1}))(p_{t+1} - \Psi_{t+1}),$$

and the adjustment is proportional to $(p_{t+1} - \Psi_{t+1})$, as though expecting the next realization to compensate for the discrepancy between p_{t+1} and the past empirical frequency. This is a form of negative correlation with past realizations as in the *gambler's fallacy*.

¹³We considered naming these biases underreaction and overreaction respectively, because attaching too much weight to the prior (as in positive prior-bias) presumably means that in a sense too little weight is attached to data (and similarly for the other axiom). However, the term underreaction suggests low sensitivity of the posterior to the signal s_{t+1} , which need not be the case in (9) unless α_{t+1} and λ_{t+1} do not depend on s_{t+1} . See Section 3.2 for more on underreaction and overreaction.

¹⁴When $\lambda_{t+1} < 0$ in (10), q_{t+1} is well-defined as a probability measure only under special conditions; for example, it suffices that $-\lambda_{t+1}/(1 - \lambda_{t+1}) \leq \min_{s_{t+2}} p_{t+1}(s_{t+2} | s_{t+1})$.

In each case the agent is assumed to suffer from the indicated fallacy at all times and histories. However, it is intuitive that she may move from one fallacy to another depending on the sample history. Thus one would like a theory that explains which fallacy applies at each history. Our framework gives this task a concrete form: one must explain ‘only’ how the weights λ_{t+1} vary with history.

Because she uses the empirical frequency measure to summarize past observations, the temptation facing an agent satisfying sample-bias depends equally on all past observations, although it might seem more plausible that more recent observations have a greater impact on temptation. This can be accommodated if Ψ_{t+1} is redefined as a weighted empirical frequency measure

$$\Psi_{t+1}(\cdot) = \sum_1^{t+1} w_{\tau,t+1} \delta_{s_\tau}(\cdot).$$

Here $\delta_{s_\tau}(\cdot)$ is the Dirac measure on the observation at time τ and $w_{\tau,t+1} \geq 0$ are weights; the special case $w_{\tau,t+1} = 1/(t+1)$ for all τ yields the earlier model. An agent who is influenced only by the most recent observation is captured by the law of motion

$$m_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1}))p_{t+1} + \lambda_{t+1}(1 - \alpha_{t+1})\delta_{s_{t+1}}.$$

If $\lambda_{t+1} < 0$, the resulting model admits an interpretation (in terms of sampling without replacement from changing urns) analogous to that offered by Rabin (2002) for his model of the law of small numbers.

3.2 Learning about parameters

This section specializes our model so as to capture the case where the data generating process is unknown up to a parameter $\theta \in \Theta$. In the benchmark Bayesian model, for any sequence of signals of length T , time t beliefs have the form

$$P_t(\cdot) = \int_{\Theta} \otimes_{t+1}^T \ell(\cdot | \theta) d\mu_t, \tag{11}$$

where $\ell(\cdot | \theta)$ is a likelihood function (measure on S), μ_0 represents prior beliefs on Θ , and μ_t denotes Bayesian posterior beliefs about the parameter at time t and after observations s_1^t . The de Finetti Theorem shows that beliefs admit such a representation if and only if P_0 is exchangeable. We describe (without axiomatic foundations) a generalization of (11) that accommodates non-Bayesian updating.

To accommodate parameters, adopt a suitable specification for (p_t, q_t) , taking (α_t) , δ , and u as given. We fix (Θ, ℓ, μ_0) and suppose for now that we are also given a process (ν_t) , where each ν_t is a probability measure on Θ . (The σ -algebra associated with Θ is suppressed.) The prior μ_0 on Θ induces time 0 beliefs about S_1 given by

$$p_0(\cdot) = m_0(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d\mu_0.$$

Proceed by induction: suppose that μ_t has been constructed and define μ_{t+1} by

$$\mu_{t+1} = \alpha_{t+1} BU(\mu_t; s_{t+1}) + (1 - \alpha_{t+1})v_{t+1}, \tag{12}$$

where $BU(\mu_t; s_{t+1})(\cdot)$ is the Bayesian update of μ_t . This equation constitutes the *law of motion* for beliefs about parameters. Finally, define (p_{t+1}, q_{t+1}) by

$$p_{t+1}(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d(BU(\mu_t; s_{t+1}))$$

$$q_{t+1}(\cdot) = \int_{\Theta} \ell(\cdot | \theta) dv_{t+1}.$$

This completes the specification of the model for any given process (v_t) .

Notice that

$$m_{t+1}(\cdot) = \alpha_{t+1} p_{t+1} + (1 - \alpha_{t+1}) q_{t+1} = \int_{\Theta} \ell(\cdot | \theta) d\mu_{t+1}.$$

In light of the discussion in Section 2.3, preferences at $t + 1$ are based on the beliefs about parameters represented by μ_{t+1} . If $\alpha_{t+1} \equiv 1$, then (μ_t) is the process of Bayesian posteriors and the above collapses to the exchangeable model (11).¹⁵ More generally, differences from the Bayesian model depend on (v_t) , examples of which are given next.¹⁶

Prior-Bias with Parameters Consider first the case where

$$v_{t+1} = (1 - \lambda_{t+1}) BU(\mu_t; s_{t+1}) + \lambda_{t+1} \mu_t,$$

where $\lambda_{t+1} \leq 1$. This is readily seen to imply (8) and hence prior-bias; the bias is positive or negative according to the sign of the λ 's. Posterior beliefs about parameters satisfy the law of motion

$$\mu_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1})) BU(\mu_t; s_{t+1}) + \lambda_{t+1}(1 - \alpha_{t+1}) \mu_t.$$

The latter equation reveals something of how the inferences of an agent with prior-bias differ from those of a Bayesian updater. We find that (assuming $\alpha_{t+1} \neq 1$)

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} < \frac{\ell(s_{t+1} | \theta)}{\ell(s_{t+1} | \theta')} \frac{\mu_t(\theta)}{\mu_t(\theta')} \quad \text{if and only if} \quad \lambda_{t+1} \ell(s_{t+1} | \theta') < \lambda_{t+1} \ell(s_{t+1} | \theta). \tag{13}$$

For a concrete example, consider coin tossing, with $S = \{H, T\}$, $\Theta \subset (0, 1)$, and $\ell(H | \theta) = \theta$ and consider beliefs after a string of H 's. If there is positive prior-bias (positive λ 's), then repeated application of (13) establishes that the agent underinfers in the sense that

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} < \frac{\mu_{t+1}^B(\theta)}{\mu_{t+1}^B(\theta')}, \quad \theta > \theta',$$

where μ_{t+1}^B is the posterior of a Bayesian who has the same prior at time 0. Similarly, negative prior-bias leads to overinference.

¹⁵Recall that α_0 is not defined for the representation.

¹⁶One general point is that, in contrast to the exchangeable Bayesian model, μ_{t+1} depends not only on the set of past observations, but also on the order in which they were realized.

Sample-Bias with Parameters Learning about parameters is consistent also with sample-bias. Take as primitive a process (ψ_{t+1}) of probability measures on Θ that provides a representation for empirical frequency measures Ψ_{t+1} of the form

$$\Psi_{t+1} = \int \ell(\cdot | \theta) d\psi_{t+1}(\theta). \tag{14}$$

Let μ_0 be given and define μ_{t+1} and ν_{t+1} inductively for $t \geq 0$ by (12) and

$$\nu_{t+1} = (1 - \lambda_{t+1})BU(\mu_t, s_{t+1}) + \lambda_{t+1}\psi_{t+1},$$

for $\lambda_{t+1} \leq 1$. Then one obtains a special case of sample-bias; the bias is positive or negative according to the sign of the λ 's. The implied law of motion for posteriors is

$$\mu_{t+1} = (1 - \lambda_{t+1}(1 - \alpha_{t+1}))BU(\mu_t; s_{t+1}) + \lambda_{t+1}(1 - \alpha_{t+1})\psi_{t+1}. \tag{15}$$

To illustrate, suppose that $S = \{s^1, \dots, s^K\}$ and $\ell(s^k | \theta) = \theta_k$ for each $\theta = (\theta_1, \dots, \theta_K)$ in Θ , the interior of the K -simplex. Then one can ensure (14) by taking ψ_0 to be a suitable noninformative prior; subsequently, Bayesian updating leads to the desired process (ψ_{t+1}) . For example, the improper Dirichlet prior density

$$\frac{d\psi_0(\theta)}{\times_{k=1}^K d\theta_k} \propto \times_{k=1}^K \theta_k^{-1}$$

yields the Dirichlet posterior with parameter vector $(n_t(s^1), \dots, n_t(s^K))$, where $n_t(s^k)$ equals the number of realizations of s^k in the first t periods; that is,

$$\frac{d\psi_t(\theta)}{\times_{k=1}^K d\theta_k} \propto \times_{k=1}^K \theta_k^{n_t(s^k)-1}. \tag{16}$$

By the property of the Dirichlet distribution,

$$\int \ell(s^k | \theta) d\psi_t(\theta) = \int \theta_k d\psi_t(\theta) = \frac{n_k(t)}{t},$$

the empirical frequency of s^k , as required by (14).

Finally, we find from (15) and (16) that (assuming $\alpha_{t+1} \neq 0$)

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} > \frac{\ell(s_{t+1} | \theta) \mu_t(\theta)}{\ell(s_{t+1} | \theta') \mu_t(\theta')} \quad \text{if and only if} \quad \lambda_{t+1} \frac{\psi_t(\theta)}{\psi_t(\theta')} > \lambda_{t+1} \frac{\mu_t(\theta)}{\mu_t(\theta')}. \tag{17}$$

Suppose that all λ_{t+1} 's are negative (negative sample-bias) and consider the coin-tossing example. As above, we denote by (μ_t^B) the Bayesian process of posteriors with initial prior $\mu_0^B = \mu_0$. Then it follows from repeated application of (16) and (17) that

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} > \frac{\mu_{t+1}^B(\theta)}{\mu_{t+1}^B(\theta')},$$

if $s_1^{t+1} = (H, \dots, H)$, $|\theta - \frac{1}{2}| > |\theta' - \frac{1}{2}|$, and the common initial prior μ_0 is uniform.¹⁷ After seeing a string of H 's the agent described herein exaggerates (relative to a Bayesian) the relative likelihoods of extremely biased coins. If instead we consider a point at which the history s_1^{t+1} has an equal number of realizations of T and H , then

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(1-\theta)} > \frac{\theta}{1-\theta} \frac{\mu_t(\theta)}{\mu_t(1-\theta)} = \frac{BU(\mu_t, H)(\theta)}{BU(\mu_t, H)(1-\theta)}$$

for any θ such that $\mu_t(\theta) > \mu_t(1-\theta)$. If there have been more realizations of H , then the preceding displayed inequality holds if

$$\left(\frac{\theta}{1-\theta}\right)^{n_{t+1}(H)-n_{t+1}(T)} < \frac{\mu_t(\theta)}{\mu_t(1-\theta)},$$

for example if $\theta < \frac{1}{2}$ and $\mu_t(\theta) \geq \mu_t(1-\theta)$. Note that the bias in this case is towards coins that are less biased ($\theta < \frac{1}{2}$). The opposite biases occur in the case of positive sample bias.

4. AXIOMATIC FOUNDATIONS

4.1 The general model

In what follows, states s vary over S , consumption c varies over C , and unless otherwise specified, time t varies over $0, 1, \dots$. A generic element of $C_t \times \mathcal{C}_t$ is $f_t = (c_t, F_t)$; t -subscripts are dropped where there is no risk of confusion. Denote by $[G_{-s_{t+1}}, M]$ the contingent menu in \mathcal{C}_t that yields $G(s'_{t+1})$ if $s'_{t+1} \neq s_{t+1}$ and M otherwise. The menu M is identified with the constant contingent menu that delivers M in all states.

The first two axioms are standard.

AXIOM 1 (Order). \succeq_t is complete and transitive.

AXIOM 2 (Continuity). Both $\{f \in C_t \times \mathcal{C}_t : f \succeq_t g\}$ and $\{f \in C_t \times \mathcal{C}_t : g \succeq_t f\}$ are closed.

In [Appendix A](#), we describe a way to mix any two elements in $C_t \times \mathcal{C}_t$. Thus we can state the Independence axiom appropriate for our setting.

AXIOM 3 (Independence). For every $0 < \lambda \leq 1$ and all f and g in $C_t \times \mathcal{C}_t$,

$$f \succeq_t g \iff \lambda f + (1-\lambda)h \succeq_t \lambda g + (1-\lambda)h.$$

The intuition for Independence is similar to that provided in [Epstein \(2006\)](#) for a three-period setting, and thus we do not elaborate here.

Given two contingent menus F and G in \mathcal{C}_t , define their union statewise, that is,

$$(F \cup G)(s) = F(s) \cup G(s).$$

The counterpart of GP's central axiom is the following.

¹⁷More generally, the latter two conditions can be replaced by $\theta'(1-\theta')/(\theta(1-\theta)) > \mu_0(\theta)/\mu_0(\theta')$.

AXIOM 4 (Set-Betweenness). For all states s_{t+1} , consumption $c \in C_t$, and all F and G in \mathcal{C}_t such that $G(s'_{t+1}) = F(s'_{t+1})$ for all $s'_{t+1} \neq s_{t+1}$,

$$(c, F) \succeq_t (c, G) \implies (c, F) \succeq_t (c, F \cup G) \succeq_t (c, G).$$

Since immediate consumption and the outcome in states other than s_{t+1} is the same in all the above rankings, the axiom is essentially a statement about how the agent feels about receiving the menus $F(s_{t+1})$, $G(s_{t+1})$, or $F(s_{t+1}) \cup G(s_{t+1})$ conditional on s_{t+1} . As a statement about the ranking of menus, Set-Betweenness may be understood as the behavioral manifestation of temptation and self-control. GP show this in their setting and Epstein (2006) adapts their interpretation to the domain of (three-period) contingent menus. The ranking of (c, F) and $(c, F \cup G)$ reveals anticipation of temptation: the strict preference

$$(c, F) \succ_t (c, F \cup G) \tag{18}$$

suggests that the decision-maker prefers that some elements of $G(s_{t+1})$ not be available as an option conditional on s_{t+1} , and presumably this preference for commitment reveals that she anticipates being tempted by some element of $G(s_{t+1})$ when choosing from the menu $F(s_{t+1}) \cup G(s_{t+1})$ conditional on s_{t+1} . For perspective, note that temptations do not exist for a standard decision-maker who evaluates a menu by its best element. In particular, she does not exhibit a preference for commitment and satisfies the stronger axiom

$$F \succeq_t G \implies F \sim_t F \cup G$$

for all F and G that agree in all but one state s . Following Kreps (1988, Ch. 13), we call this axiom *strategic rationality*.

Set-Betweenness allows us to infer the agent's anticipated time $t + 1$ choices from menus, for example, whether she expects to succumb to temptation or to exert self-control. To illustrate, suppose that $F = [H_{-s_{t+1}}, \{f\}]$ and $G = [H_{-s_{t+1}}, \{g\}]$ and also that the decision-maker exhibits the preference

$$(c, [H_{-s_{t+1}}, \{f\}]) \succ_t (c, [H_{-s_{t+1}}, \{g\}]). \tag{19}$$

This ranking suggests that from the ex ante perspective of period t , she prefers to end up with f rather than with g conditional on s_{t+1} , and in particular, that she prefers f to be chosen from $\{f, g\}$ conditional on s_{t+1} . Whether she anticipates f actually being chosen from $\{f, g\}$ is then revealed by her ranking of $(c, [H_{-s_{t+1}}, \{f, g\}])$ and $(c, [H_{-s_{t+1}}, \{g\}])$. For instance, if

$$(c, [H_{-s_{t+1}}, \{f, g\}]) \succ_t (c, [H_{-s_{t+1}}, \{g\}]), \tag{20}$$

then she has a strict preference for f being available ex post, which reveals that she anticipates choosing f from $\{f, g\}$ at $t + 1$. On the other hand, if she is indifferent to f being available ex post, that is,

$$(c, [H_{-s_{t+1}}, \{f, g\}]) \sim_t (c, [H_{-s_{t+1}}, \{g\}]), \tag{21}$$

then she anticipates a weak preference at $t + 1$ for choosing g from $\{f, g\}$. To see this, observe that given (19), (21) implies (18), which in turn implies that g is tempting. Thus,

the indifference in (21) implies that she expects either to submit to g , or to be indifferent between submitting to g and resisting it. That is, she anticipates a weak preference for g at $t + 1$.

Discussion of (20)–(21) revolved around what the decision-maker *anticipates* at time t about her choices at time $t + 1$. The next axiom connects her time t expectations regarding future behavior and her *actual* future behavior.

AXIOM 5 (Sophistication). *If $(c, [G_{-s_{t+1}}, \{f\}]) \succ_t (c, [G_{-s_{t+1}}, \{g\}])$, then*

$$(c, [G_{-s_{t+1}}, \{f, g\}]) \succ_t (c, [G_{-s_{t+1}}, \{g\}]) \iff f \succ_{t+1} g,$$

where \succeq_t and \succeq_{t+1} correspond to histories (s_1, \dots, s_t) and $(s_1, \dots, s_t, s_{t+1})$ respectively.

The axiom states that she is sophisticated in that her expectations are correct (at least for anticipated choices out of binary menus $\{f, g\}$). To see this, start by taking f, g such that in period t she would prefer to commit to f rather than g conditionally on s_{t+1} (as in the hypothesis). As in the earlier discussion, this relationship between f and g allows us to deduce her expected $t + 1$ choice out of $\{f, g\}$ from her \succeq_t -ranking of $(c, [G_{-s_{t+1}}, \{f, g\}])$ and $(c, [G_{-s_{t+1}}, \{g\}])$. Her actual choice out of $\{f, g\}$ is given by her \succeq_{t+1} -ranking of f and g . The axiom states that the decision-maker expects to choose f at $t + 1$ if and only if she in fact chooses f at $t + 1$.

Some axioms below involve the evaluation of streams of lotteries (or lottery streams), and it is convenient to introduce relevant notation at this point. Any risky consumption stream for the time period $[t + 1, \infty)$, that is, where a unique (independent of states) consumption level c_τ is prescribed for each $\tau \geq t + 1$, may be identified with an element of $C_{t+1} \times C_{t+2} \times \dots$. Denote by \mathcal{L}_{t+1} the subset of all such risky consumption streams; a generic element is $\ell = (\ell_\tau)_{\tau=t+1}^\infty$.

In order to obtain meaningful probabilities, a form of state independence is needed.

AXIOM 6 (State Independence). *For all s_{t+1} , contingent menus F in \mathcal{C}_{t+1} and $\ell', \ell \in \mathcal{L}_{t+1}$,*

$$(c, \{\ell'\}) \succeq_t (c, \{\ell\}) \iff (c, [F_{-s_{t+1}}, \{\ell'\}]) \succeq_t (c, [F_{-s_{t+1}}, \{\ell\}]).$$

The axiom states that the ranking of the lottery streams ℓ' and ℓ received unconditionally does not change if they are received conditionally on any specific s_{t+1} obtaining. Thus time preferences and risk attitudes are not state-dependent.

In our model, temptation arises only because of a change in beliefs. This is reflected in the next axiom.¹⁸

AXIOM 7 (Restricted Strategic Rationality (RSR)). *For all states s_{t+1}, s_{t+2} , consumption $c, c' \in C$, and contingent menus $F \in \mathcal{C}_t$ and $H, H' \in \mathcal{C}_{t+1}$ such that $H(s'_{t+2}) = H'(s'_{t+2})$ for all $s'_{t+2} \neq s_{t+2}$, if*

$$(c', [F_{-s_{t+1}}, \{(c, H')\}]) \succeq_t (c', [F_{-s_{t+1}}, \{(c, H)\}]), \quad (22)$$

¹⁸As in Sophistication, the preferences \succeq_t and \succeq_{t+1} correspond to histories (s_1, \dots, s_t) and $(s_1, \dots, s_t, s_{t+1})$ respectively.

then

$$(c', [F_{-s_{t+1}}, \{(c, H')\}]) \sim_t (c', [F_{-s_{t+1}}, \{(c, H'), (c, H)\}]) \tag{23}$$

$$(c, H') \succeq_{t+1} (c, H). \tag{24}$$

Suppose that, on observing s_{t+1} , the agent at $t + 1$ has to choose from the menu $\{(c, H'), (c, H)\}$ where $H'(s'_{t+2}) = H(s'_{t+2})$ for all $s'_{t+2} \neq s_{t+2}$ for some s_{t+2} . Since H' and H differ only in the single state s_{t+2} , their ranking does not depend on beliefs over S_{t+2} : there are no trade-offs across states that must be made. Consequently, there is no temptation when choosing out of $\{(c, H'), (c, H)\}$, and, therefore, conditional on any s_{t+1} , the agent never exhibits a preference for commitment. In particular, her preference \succeq_t satisfies a form of strategic rationality. This is the content of the implication '(22) \implies (23)'. The implication '(22) \implies (24)' is another expression of the absence of temptation: if the $t + 1$ choice between the prospects (c, H') and (c, H) is not subject to temptation, then there is no reason for her $t + 1$ perspective to deviate from her prior, time t perspective regarding the two prospects. The latter perspective is revealed by (22), the agent's time t preference for committing to (c, H') versus (c, H) conditionally on s_{t+1} .

The final axiom places structure on the agent's preferences over lottery streams.

AXIOM 8 (Risk Preference). *There exist $0 < \delta < 1$ and $u : C \rightarrow \mathbb{R}^1$ nonconstant, linear and continuous, such that, for each ℓ' and ℓ in \mathcal{L}_{t+1} ,*

$$\ell' \succeq_t \ell \iff \sum_{t+1}^{\infty} \delta^{\tau-(t+1)} u(\ell'_\tau) \geq \sum_{t+1}^{\infty} \delta^{\tau-(t+1)} u(\ell_\tau).$$

The axiomatic characterization of the utility function over streams of lotteries appearing in (8) is well known (see Epstein 1983, for example). Because time and risk preferences are not our primary focus, we content ourselves with the statement of the above unorthodox 'axiom.'

Say that $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{1 \leq t \leq T})$ represents (\succeq_t) if \succeq_0 is represented by $\mathcal{U}_0(\cdot)$ and, for each $t > 0$, \succeq_t is represented by $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$, where these functions are defined in (5)–(7) and where u, δ, p_0 , and $(\alpha_t, p_t, q_t)_{t \geq 1}$ satisfy the properties stated there. For any $c \in C_{t+1}$ and $M \subset \mathcal{C}_{t+1}$, write (c, M) instead of $\{c\} \times M \in \mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1})$.

THEOREM 1. *If the process of preferences (\succeq_t) satisfies Axioms 1–8, then there exists some $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{t \geq 1})$ representing (\succeq_t) .*

Conversely, suppose that

$$(1 + 2 \sup_{t, s_1^t} |\alpha_t^{-1} - 1|) \delta < 1. \tag{25}$$

Then equations (5)–(7) admit a unique solution (U_t) , where $U_t(\cdot, s_1^t) : \mathcal{M}(C \times \mathcal{C}) \rightarrow \mathbb{R}^1$ is continuous and uniformly bounded in the sense that

$$\|(U_t)\| \equiv \sup_{t, s_1^t, M} |U_t(M; s_1^t)| < \infty.$$

Define $\mathcal{U}_t(\cdot, s_t^t)$ and $\mathcal{V}_t(\cdot, s_t^t)$ by (5)–(6) and let \succeq_0 be represented by $\mathcal{U}_0(\cdot)$, and, for each $t > 0$, let \succeq_t be represented by $\alpha_t \mathcal{U}_t(\cdot) + (1 - \alpha_t) \mathcal{V}_t(\cdot)$. Then (\succeq_t) satisfies Axioms (1)–(8).

Finally, if $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{t \geq 1})$ and $(\delta', u', p'_0, (\alpha'_t, p'_t, q'_t)_{t \geq 1})$ both represent (\succeq_t) , then $\delta' = \delta$, $u' = a u + b$ for some $a > 0$, and

$$p'_0 = p_0, \alpha'_t p'_t + (1 - \alpha'_t) q'_t = \alpha_t p_t + (1 - \alpha_t) q_t \text{ for } t > 0.$$

If t and s_{t+1} are such that

$$(F_{-s_{t+1}}, (c, M')) \succ_t (F_{-s_{t+1}}, (c, M' \cup M)) \tag{26}$$

for some $c \in C_{t+1}$ and $M', M \subset \mathcal{C}_{t+1}$, then

$$(\alpha'_{t+1}(s_{t+1}), q'_{t+1}(\cdot | s_{t+1})) = (\alpha_{t+1}(s_{t+1}), q_{t+1}(\cdot | s_{t+1})). \tag{27}$$

The restriction (25) implies that the recursion (7) defines a contraction mapping which then yields a unique solution. The second part of the theorem deals with uniqueness. Absolute uniqueness of all components is not to be expected. For example, if $\alpha_{t+1}(s_{t+1}) = 0$, then every measure $q_{t+1}(\cdot | s_{t+1})$ leads to the same s_{t+1} -conditional preference; similarly, if $q_{t+1}(\cdot | s_{t+1}) = p_{t+1}(\cdot | s_{t+1})$, then $\alpha_{t+1}(s_{t+1})$ is of no consequence and hence indeterminate. These degenerate cases constitute precisely the circumstances under which s_{t+1} -conditional preference is strategically rational, which is what is excluded by condition (26). Once strategic rationality is excluded, the strong uniqueness property in (27) obtains.

4.2 Foundations for prior-bias and sample-bias

The specializations prior-bias and sample-bias introduced in Section 3 are characterized here (we use upper case names for the axioms corresponding to each bias).

It is convenient to define the preference $\succeq_t|_{s_{t+1}}$ on $\mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1})$ by

$$M \succeq_t|_{s_{t+1}} M' \iff (c, [H_{-s_{t+1}}, M]) \succeq_t (c, [H_{-s_{t+1}}, M'])$$

for some $(c, H) \in C_t \times \mathcal{C}_t$; the additive separability of the representation ensures that the choice of (c, H) is irrelevant. For any $(c_{t+1}, F_{t+1}) \in C_{t+1} \times \mathcal{C}_{t+1}$, denote by $(c, \{(c_{t+1}, F_{t+1})\})$ the alternative that yields immediate consumption c and a contingent menu that commits the agent to (c_{t+1}, F_{t+1}) in every state s_{t+1} . Evidently, the evaluation of any such prospect reflects marginal beliefs about S_{t+2} held at time t , that is, the agent's period t prior on S_{t+2} . Say that s_{t+1} is a *neutral signal* if, for all $c_t, c_{t+1} \in C$ and $F_{t+1}, G_{t+1} \in \mathcal{C}_{t+1}$,

$$\{(c_{t+1}, F_{t+1})\} \succeq_t|_{s_{t+1}} \{(c_{t+1}, G_{t+1})\} \iff (c_t, \{(c_{t+1}, F_{t+1})\}) \succeq_t (c_t, \{(c_{t+1}, G_{t+1})\}).$$

Given our representation, s_{t+1} is a neutral signal if and only if $p_{t+1}(s_{t+2} | s_{t+1}) = \int p_{t+1}(s_{t+2} | s'_{t+1}) dm_t(s'_{t+1})$ for all s_{t+2} .

AXIOM 9 (Prior-Bias). Let $s_{t+1} \in S_{t+1}$ and suppose that for $c \in C_{t+1}$ and $F_{t+1}, G_{t+1} \in \mathcal{C}_{t+1}$,

$$\{(c, F_{t+1})\} \succ_{t|s_{t+1}} \{(c, G_{t+1})\}. \tag{28}$$

If either s_{t+1} is a neutral signal, or if, for some $c_t \in C_t$,

$$(c_t, \{(c, F_{t+1})\}) \sim_t (c_t, \{(c, G_{t+1})\}), \tag{29}$$

then

$$\{(c, F_{t+1})\} \sim_{t|s_{t+1}} \{(c, F_{t+1}), (c, G_{t+1})\}. \tag{30}$$

To interpret the axiom, we suppress the fixed consumption c_t and c_{t+1} (and do the same for interpretations in the sequel). Condition (28) states that at time t , the agent strictly prefers to commit to F rather than to G conditionally on s_{t+1} . There are two situations in which she would not be tempted by G conditionally on s_{t+1} at time $t + 1$ (and thus not exhibit a preference for commitment (30)). The first is when s_{t+1} is a neutral signal, and thus does not lead to any updating of the prior. The second is when she is indifferent between F and G if they are received unconditionally (29), that is, if prior beliefs about S_{t+2} make both look equally attractive. That the presence of temptation conditionally on s_{t+1} depends not only on how F and G are ranked conditionally but also on how attractive they were prior to the realization of s_{t+1} , indicates excessive influence of prior beliefs at the updating stage (time $t + 1$).

Prior-Bias begs the question what happens to temptation if the indifference in (29) is not satisfied. We consider two alternative strengthenings of the axiom that provide different answers.

Label by *Positive Prior-Bias* the axiom obtained when (29) is replaced by

$$(c', \{(c, F)\}) \succeq_t (c', \{(c, G)\}).$$

This implies that G is tempting conditionally on s_{t+1} only if it was more attractive according to (time t) prior beliefs about S_{t+2} . An alternative, labeled *Negative Prior-Bias*, is the axiom obtained when (29) is replaced by

$$(c', \{(c, F)\}) \preceq_t (c', \{(c, G)\}).$$

In this case, G is preferred at time t , but the signal s_{t+1} reverses the ranking in favor of F . Thus s_{t+1} is a strong positive signal for F . The agent is greatly influenced by signals. Thus she is not tempted by G after seeing s_{t+1} .

COROLLARY 1. Suppose that (\succeq_t) has a representation $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{t \geq 1})$. Then (\succeq_t) satisfies Prior-Bias if and only if it admits a representation satisfying (8) for some adapted process (λ_t) with $\lambda_{t+1} \leq 1$. Further, (\succeq_t) satisfies (i) Positive Prior-Bias or (ii) Negative Prior-Bias if and only if (8) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$.¹⁹

¹⁹The proofs of this corollary and of the one to follow are similar to the proof of Epstein (2006, Corollary 3.5), and thus we do not include them.

Sample-bias can be characterized along the same lines. We need some additional notation: denote by Ψ_{t+1} the empirical frequency measure on S given the history s_1^{t+1} . For any G in \mathcal{C}_{t+1} , $G(s_{t+2})$ is a subset of $C_{t+2} \times \mathcal{C}_{t+2}$ and so is the mixture $\int G(s'_{t+2})d\Psi_{t+1}$. Consider the contingent menu in \mathcal{C}_{t+1} , denoted $\int Gd\Psi_{t+1}$, that assigns $\int G(s'_{t+2})s d\Psi_{t+1}$ to every s_{t+2} . Then $(c_{t+1}, \int Gd\Psi_{t+1})$ denotes an alternative that yields the obvious singleton menu.

The axioms to follow parallel the trio of axioms just stated. One difference is that the contingent menus F and G appearing in these axioms are assumed, for reasons given below, to lie in $\mathcal{C}_{t+1}^{c,+1} \subset \mathcal{C}_{t+1}$. Thus F and G provide perfect commitment and are such that all relevant uncertainty is resolved by $t + 2$. In this setting, say that s_{t+1} is a *neutral signal* if, for all $c \in C_{t+1}$ and $F, G \in \mathcal{C}_{t+1}^{c,+1}$,

$$\{(c, F)\} \succeq_{t|s_{t+1}} \{(c, G)\} \iff \left\{ \left(c, \int F d\Psi_{t+1} \right) \right\} \succeq_{t|s_{t+1}} \left\{ \left(c, \int G d\Psi_{t+1} \right) \right\}. \tag{31}$$

The right-hand side can be interpreted as saying that *the sample s_1^{t+1} makes F look more attractive than G* : F delivers $F(s_{t+2})$ in state s_{t+2} and s_{t+2} appears with frequency $\Psi_{t+1}(s_{t+2})$ in the sample. Thus ‘on average’, F yields $\int Fd\Psi_{t+1}$. But the agent is indifferent between F and its average because she satisfies Independence. Thus the right-hand side in (31) implies that, under Ψ_{t+1} , the average for F is better than that of G . Thus for a neutral signal s_{t+1} , F is more attractive than G under commitment (conditional on s_{t+1}) if and only if F is more attractive than G on average under the sample history. Given our representation, s_{t+1} is a neutral signal if and only if $p_{t+1}(s_{t+2} | s_{t+1}) = \Psi_{t+1}(s_{t+2} | s_{t+1})$ for all s_{t+2} .

AXIOM 10 (Sample-Bias). For $s_{t+1} \in S_{t+1}$, $c \in C_{t+1}$ and F, G in $\mathcal{C}_{t+1}^{c,+1}$ such that

$$\{(c, F)\} \succ_{t|s_{t+1}} \{(c, G)\},$$

if either s_{t+1} is a neutral signal, or if, for some $c_t \in C_t$,

$$\left\{ \left(c, \int F d\Psi_{t+1} \right) \right\} \sim_{t|s_{t+1}} \left\{ \left(c, \int G d\Psi_{t+1} \right) \right\}, \tag{32}$$

then

$$\{(c, F)\} \sim_{t|s_{t+1}} \{(c, F), (c, G)\}.$$

The next two axioms provide alternative strengthenings of Sample-Bias. Label by *Positive Sample-Bias* the axiom obtained if (32) is replaced by

$$\left\{ \left(c, \int F d\Psi_{t+1} \right) \right\} \succeq_{t|s_{t+1}} \left\{ \left(c, \int G d\Psi_{t+1} \right) \right\}.$$

Similarly, ‘define’ *Negative Sample-Bias* by using the hypothesis

$$\left\{ \left(c, \int F d\Psi_{t+1} \right) \right\} \preceq_{t|s_{t+1}} \left\{ \left(c, \int G d\Psi_{t+1} \right) \right\}.$$

We now interpret Positive Sample-Bias; the other interpretations are similar. The axiom asserts that if commitment to F is preferred (conditionally on s_{t+1}) to commitment to G , and if the sample makes F look more attractive than G , or if s_{t+1} is neutral, then G is not tempting conditionally. The fact that the sample may influence temptation after realization of s_{t+1} , above and beyond its role in the conditional ranking, reveals the excessive influence of the sample at the updating stage. The influence is ‘positive’ because G can be tempting conditionally only if it was more attractive according to the sample history.

The preceding intuition, specifically the indifference between F and $\int F d\Psi_{t+1}$ posited when interpreting (31), relies on F lying in $\mathcal{C}_{t+1}^{c,+1}$. That is because as s_{t+2} varies, not only does $F(s_{t+2})$ vary but so also does the information upon which the agent bases evaluation of the menu $F(s_{t+2})$. Independence implies indifference to the former variation but not to the latter. For F in $\mathcal{C}_{t+1}^{c,+1}$, however, information is irrelevant because all uncertainty is resolved once s_{t+2} is realized.

COROLLARY 2. *Suppose that (\succeq_t) has a representation $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{t \geq 1})$. Then (\succeq_t) satisfies Sample-Bias if and only if it admits a representation satisfying (10) for some adapted process (λ_t) with $\lambda_{t+1} \leq 1$. Further, (\succeq_t) satisfies (i) Positive Sample-Bias or (ii) Negative Sample-Bias if and only if (10) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$.*

APPENDIX

A. CONTINGENT MENUS

The construction of the space of contingent menus is analogous to familiar constructions of type spaces (Mertens and Zamir 1985 and Brandenburger and Dekel 1993), and to related constructions in Epstein and Wang (1996) and Gul and Pesendorfer (2004). The difficulty arises from a problem of infinite regress. In the context of type spaces, the solution is to employ suitable hierarchies of spaces of probability measures. Here and in the other studies cited, hierarchies of alternative topological spaces are used. The technical details are now well understood and thus we omit a formal proof for the theorem that follows.²⁰ The properties of the space of contingent menus spelled out in the theorem are invoked in proving our main representation result **Theorem 1**. Readers who are not interested in that proof and who are willing to accept the intuitive description of contingent menus provided in the discussion leading to (3) may skip this section of the appendix entirely.

Define the following spaces:

$$D_1 = [\mathcal{M}(C \times C^\infty)]^S$$

$$D_t = [\mathcal{M}(C \times D_{t-1})]^S \quad \text{for } t > 1.$$

²⁰When S is a singleton, a contingent menu is simply a menu, hence a closed subset, and the proof is a corollary of Epstein and Wang (1996, Theorem 6.1). See also Epstein and Peters (1999, Appendix B), which deals with hierarchies of upper-semicontinuous functions taking values in $[0, 1]$; the indicator function of a closed set is such a function, hence the relevance to hierarchies of closed sets.

For interpretation, G in D_1 yields the set $G(s)$ of consumption streams if s is realized at $t = 1$. Thus think of G as a contingent menu for which there is no uncertainty and no flexibility (in the sense of nonsingleton menus) after time 1. Similarly, G in D_t can be thought of as a contingent menu for which there is no uncertainty or flexibility after time t .

Each D_t is compact metric. In addition, there is a natural mixing operation on each D_t : given any space X where mixtures $\lambda x + (1 - \lambda)y$ are well defined, mix elements of $\mathcal{M}(X)$ by

$$\lambda M + (1 - \lambda)N = \{\lambda x + (1 - \lambda)y : x \in M, y \in N\}.$$

Mixtures are defined in the obvious way on $X = C^\infty$. On D_1 define $\lambda G' + (1 - \lambda)G$ by

$$(\lambda G' + (1 - \lambda)G)(s) = \lambda G'(s) + (1 - \lambda)G(s).$$

Proceed inductively for all D_t .

THEOREM A.1. *There exists $\mathcal{C} \subset \times_1^\infty D_t$ such that*

- (i) \mathcal{C} is compact metric under the induced product topology.
- (ii) \mathcal{C} is homeomorphic to $[\mathcal{M}(C \times \mathcal{C})]^S$.
- (iii) Under a suitable identification,

$$D_{t-1} \subset D_t \subset \mathcal{C}.$$

- (iv) Let π_t be the projection map from $\times_1^\infty D_t$ into D_t . Then $\pi_t(\mathcal{C}) \subset \mathcal{C}$ and

$$\pi_t(F) \xrightarrow[t \rightarrow \infty]{} F \text{ for every } F \text{ in } \mathcal{C}.$$

- (v) Let $F' = (G'_t)$ and $F = (G_t)$ be in \mathcal{C} . Then $(\lambda G'_t + (1 - \lambda)G_t)$ is an element of \mathcal{C} , denoted $\lambda \circ F' + (1 - \lambda) \circ F$. Under the homeomorphism in (i),

$$\begin{aligned} &(\lambda \circ F' + (1 - \lambda) \circ F)(s) \\ &= \{(\lambda c' + (1 - \lambda)c, \lambda \circ H' + (1 - \lambda) \circ H) : (c', H') \in F'(s), (c, H) \in F(s)\}. \end{aligned}$$

Part (i) asserts that the topological structure of C is inherited by \mathcal{C} . Part (ii) is the homeomorphism (3) used in the text.

We note above that each G in D_t implies no uncertainty or flexibility after time t . Think of such a G as a special contingent menu in which all uncertainty and flexibility beyond t have been somehow collapsed into period t . Then (iii) and (iv) imply that the set $\cup_1^\infty D_t$ of all such special contingent menus is dense in \mathcal{C} .

Part (v) provides the mixing operation promised in Section 4. Roughly it shows that ‘ \circ ’, which is the natural mixing operation induced by $\times_1^\infty D_t$ on \mathcal{C} , is consistent with that suggested by the homeomorphism in (ii). Thus, there is no danger of confusion and in the text we have written simply $\lambda F' + (1 - \lambda)F$ rather than $\lambda \circ F' + (1 - \lambda) \circ F$.

Finally, define the spaces $\mathcal{C}^{c,+1} \subset \mathcal{C}^c \subset \mathcal{C}$ introduced in Section 2.3. First, \mathcal{C}^c is the unique subspace of \mathcal{C} satisfying

$$\mathcal{C}^c \underset{\text{homeo}}{\approx} (C \times \mathcal{C}^c)^S$$

under the homeomorphism in the theorem. (Details are as in Epstein and Wang 1996, Theorem 6.1(a).) Take $\mathcal{C}^{c,+1} = (C \times C^\infty)^S$.

B. PROOF OF MAIN REPRESENTATION RESULT

B.1 Necessity of the axioms

Denote by X the set of all processes $U = (U_t)$, where $U_t(\cdot, s_1^t): \mathcal{M}(C \times \mathcal{C}) \rightarrow \mathbb{R}^1$ is continuous and where

$$\|U\| = \|(U_t)\| \equiv \sup_{t, s_1^t, M} |U_t(M, s_1^t)| < \infty.$$

The norm $\|\cdot\|$ makes X a Banach space. Define $\Gamma: X \rightarrow X$ by

$$\begin{aligned} (\Gamma(U))_{t+1}(M_{t+1}, s_{t+1}) = & \\ & \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \frac{1}{\alpha_{t+1}} \left\{ u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) d(\alpha_{t+1} p_{t+1} + (1 - \alpha_{t+1}) q_{t+1}) \right\} \\ & - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \frac{1 - \alpha_{t+1}}{\alpha_{t+1}} \left\{ u(c'_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F'_{t+1}(s_{t+2}), s_{t+2}) d q_{t+1}(s_{t+2}) \right\}. \end{aligned}$$

Then Γ is a contraction under assumption (25) and thus has a unique fixed point (U_t) .

It is a routine matter to verify the axioms.

B.2 Preliminaries for sufficiency

For any compact metric space D endowed with a continuous mixture operation, say that a preference \succeq over $\mathcal{M}(D)$ has a (U, V) representation if the functions $U, V: D \rightarrow \mathbb{R}$ are continuous and linear, and if \succeq is represented by $W_{U,V}: \mathcal{M}(D) \rightarrow \mathbb{R}$, where

$$W_{U,V}(M) = \max_{c \in M} \{U + V\} - \max_{c' \in M} V, \quad M \in \mathcal{M}(D).$$

Say that \succeq is *strategically rational* if for all $M, M' \in \mathcal{M}(D)$,

$$M \succeq M' \implies M \sim M \cup M'.$$

LEMMA B.1. *If \succeq has a (U, V) representation with U nonconstant, then*

- (i) \succeq is strategically rational if and only if $V = aU + b$ for some $a \geq 0$. In particular, if V is nonconstant then \succeq is strategically rational if and only if $U + V = aV + b$ for some $a > 1$.

(ii) \succeq is strategically rational if and only if for all $\bar{c}, \underline{c} \in D$,

$$\{\bar{c}\} \succeq \{\underline{c}\} \implies \{\bar{c}\} \sim \{\bar{c}, \underline{c}\}. \quad (\text{B.1})$$

PROOF. (i) The argument is similar to Gul and Pesendorfer (2001, p. 1414).

(ii) Sufficiency is clear. For necessity, suppose that \succeq is not strategically rational so that, as in Gul and Pesendorfer (2001, p. 1414), U and V are nonconstant and U is not a positive affine transformation of V . Consequently, there exist $\bar{c}, \underline{c} \in D$ such that either $[U(c) > U(c') \text{ and } V(c) \leq V(c')]$, or $[U(c) \geq U(c') \text{ and } V(c) < V(c')]$. Linearity and non-constancy of U and V imply the existence of \bar{c} and \underline{c} close to c and c' , respectively, such that all inequalities are strict. Then

$$\{\bar{c}\} \succ \{\underline{c}\} \text{ and } \{\bar{c}\} \succ \{\bar{c}, \underline{c}\},$$

which violates (B.1) and yields the result. \square

LEMMA B.2. Suppose that \succeq has a (U, V) representation and that there exist \bar{c}, \underline{c} such that $\{\bar{c}, \underline{c}\} \succ \{\underline{c}\}$. Then a preference \succeq^* over D is represented by $U + V$ if and only if it satisfies the vNM axioms and the restriction

$$\text{if } \{c\} \succ \{d\}, \text{ then } \{c, d\} \succ \{d\} \iff c \succ^* d.$$

For a proof of this result, see Noor (2008, Theorem 2.2). For any state s_{t+2} , $G \in \mathcal{C}_{t+1}$, and $L \subset \mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$, define the set $Ls_{t+2}G$ of contingent menus by

$$Ls_{t+2}G = \{[G_{-s_{t+2}}, M] : M \in L\} \subset \mathcal{C}_{t+1}.$$

Define $\succeq_{t|s_{t+1}, s_{t+2}}$ on closed subsets of $\mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$ by $L' \succeq_{t|s_{t+1}, s_{t+2}} L$ if and only if

$$(c', [F_{-s_{t+1}}, (c, L's_{t+2}G)]) \succeq_t (c', [F_{-s_{t+1}}, (c, Ls_{t+2}G)]),$$

for some $c, c' \in C$, F in \mathcal{C}_t , and G in \mathcal{C}_{t+1} .

LEMMA B.3. Suppose that (\succeq_t) satisfies Axioms 1–8 and that $\succeq_{t|s_{t+1}, s_{t+2}}$ has a (U, V) representation with nonconstant U . Then $\succeq_{t|s_{t+1}, s_{t+2}}$ is strategically rational.

PROOF. By Lemma B.1(ii), we need only establish that for any $M, M' \in \mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$,

$$\{M\} \succeq_{t|s_{t+1}, t+2} \{M'\} \implies \{M\} \sim_{t|s_{t+1}, t+2} \{M, M'\}.$$

Observe that

$$\begin{aligned} & \{M\} \succeq_{t|s_{t+1}, t+2} \{M'\} \\ & \iff (c', [F_{-s_{t+1}}, \{(c, [G_{-s_{t+2}}, M])\}]) \succeq_t (c', [F_{-s_{t+1}}, \{(c, [G_{-s_{t+2}}, M'])\}]) \\ & \implies^* (c', [F_{-s_{t+1}}, \{(c, [G_{-s_{t+2}}, M])\}]) \sim_t (c', [F_{-s_{t+1}}, \{(c, [G_{-s_{t+2}}, M]), (c, [G_{-s_{t+2}}, M'])\}]) \\ & \implies \{M\} \sim_{t|s_{t+1}, t+2} \{M, M'\}, \end{aligned}$$

where the implication \implies^* is by RSR. \square

In the next lemma, \succeq_t and \succeq_{t+1} are the preferences corresponding to histories (s_1, \dots, s_t) and $(s_1, \dots, s_t, s_{t+1})$ respectively.

LEMMA B.4. *Suppose that (\succeq_t) satisfies Axioms 1–8. If $H, H' \in \mathcal{C}_{t+1}$ are such that $H(s'_{t+2}) = H'(s'_{t+2})$ for all $s'_{t+2} \neq s_{t+2}$, then for any s_{t+1}, c, c' and F ,*

$$(c, H) \succeq_{t+1} (c, H') \iff (c', [F_{-s_{t+1}}, \{(c, H)\}]) \succeq_t (c', [F_{-s_{t+1}}, \{(c, H')\}]).$$

PROOF. \Leftarrow follows from RSR. Conversely, suppose that $(c, H) \succeq_{t+1} (c, H')$ and $(c', [F_{-s_{t+1}}, \{(c, H')\}]) \succ_t (c', [F_{-s_{t+1}}, \{(c, H)\}])$. Sophistication implies

$$(c', [F_{-s_{t+1}}, \{(c, H'), (c, H)\}]) \preceq_t (c', [F_{-s_{t+1}}, \{(c, H)\}]);$$

by Set-Betweenness, this weak preference is in fact indifference. Therefore,

$$(c', [F_{-s_{t+1}}, \{(c, H')\}]) \succ_t (c', [F_{-s_{t+1}}, \{(c, H'), (c, H)\}]),$$

which contradicts RSR. □

B.3 Sufficiency of the axioms

The proof of sufficiency begins by establishing the desired representation of \succeq_0 on $C \times D_T \subset C \times \mathcal{C}$ (see [Appendix A](#)). Later the representation of \succeq_0 is extended to all of $C \times \mathcal{C}$ by letting $T \rightarrow \infty$ and exploiting the denseness indicated in [Theorem A.1](#)(iv). The desired representations for (\succeq_t) follow.

Until specified otherwise, we derive a representation for the restriction of \succeq_0 to $C \times D_T$, for given $T > 0$. The argument involves deriving, for each $0 \leq t < T$, an appropriate representation for the restriction of \succeq_t to $C_t \times D_{T-t}$. This proceeds by backward induction on t . Define $U^r : \mathcal{M}(C \times C^\infty) \rightarrow \mathbb{R}^1$ by

$$U^r(M) = \max_{\ell \in M} \sum_0^\infty \delta^\tau u(\ell_\tau),$$

where δ and u are provided by Risk Independence. We begin by showing that \succeq_{T-1} is represented on $C_{T-1} \times D_1$ by the function

$$\mathcal{W}_{T-1}(c_{T-1}, F_{T-1}) = u(c_{T-1}) + \delta \int_{S_T} U^r(F_{T-1}(s_T)) d m_{T-1}, (c_{T-1}, F_{T-1}) \in C \times D_1, \tag{B.2}$$

where $m_{T-1} \in \Delta(S_T)$ and m_{T-1} has full support.

Identify $(C \times C^\infty)^S$ with the obvious subset of $\mathcal{M}(C \times C^\infty)$. We claim that the restriction of \succeq_{T-1} to $C_{T-1} \times (C \times C^\infty)^S$ may be represented by

$$\mathcal{W}_{T-1}(c, F_{T-1}) = u_1(c) + u_2(F_{T-1}), \tag{B.3}$$

where $u_1(\cdot)$ and $u_2(\cdot)$ are continuous and linear. We argue as follows. Since $C_{T-1} \times (C \times C^\infty)^S$ is a mixture space and \succeq_{T-1} satisfies Order, Continuity, and Independence, there exists a continuous linear representation $\mathcal{W}_{T-1}(\cdot)$ of \succeq_{T-1} on $C_{T-1} \times (C \times C^\infty)^S$. By definition of the mixture operation, for any $c, c' \in C_{T-1}$ and $F, F' \in D_1$,

$$\frac{1}{2}(c, F) + \frac{1}{2}(c', F') = \frac{1}{2}(c', F) + \frac{1}{2}(c, F').$$

Thus

$$\begin{aligned} \mathcal{W}_{T-1}(\tfrac{1}{2}(c, F) + \tfrac{1}{2}(c', F')) &= \mathcal{W}_{T-1}(\tfrac{1}{2}(c', F) + \tfrac{1}{2}(c, F')) \\ \implies \tfrac{1}{2}\mathcal{W}_{T-1}(c, F) + \tfrac{1}{2}\mathcal{W}_{T-1}(c', F') &= \tfrac{1}{2}\mathcal{W}_{T-1}(c', F) + \tfrac{1}{2}\mathcal{W}_{T-1}(c, F') \\ \implies \mathcal{W}_{T-1}(c, F) &= \mathcal{W}_{T-1}(c, F') + \mathcal{W}_{T-1}(c', F) - \mathcal{W}_{T-1}(c', F') \equiv u_1(c) + u_2(F). \end{aligned}$$

Linearity and continuity of $u_1(\cdot)$ and $u_2(\cdot)$ are evident.

Next, we show that $u_1(\cdot)$ and $u_2(\cdot)$ from (B.3) are such that it is without loss of generality to set

$$\mathcal{W}_{T-1}(c, F_{T-1}) = u(c) + \delta \int_{S_T} U'(F_{T-1}(s_T)) dm_{T-1}, \quad F_{T-1} \in (C \times C^\infty)^S,$$

for some $m_{T-1} \in \Delta(S_T)$. Take any c and define \succeq on $(C \times C^\infty)^S$ by

$$F \succeq G \iff (c, F) \succeq_{T-1} (c, G). \tag{B.4}$$

We verify that \succeq satisfies the Anscombe–Aumann axioms. Order, Continuity, and Independence are immediate. By Risk Preference and nonconstancy of $u(\cdot)$, there exist $c', c'' \in C$ such that for any $\vec{c} \in C^\infty$, $(c', \vec{c}) \not\sim (c'', \vec{c})$, and thus \succeq satisfies the Anscombe–Aumann nondegeneracy condition. State Independence applied twice yields $(F_{-s_T}, c') \succeq (F_{-s_T}, c'') \implies (F_{-s'_T}, c') \succeq (F_{-s'_T}, c'')$ for all $c', c'' \in C$ and $s_T, s'_T \in S_T$. Thus there exists $m_{T-1} \in \Delta(S_T)$ and $v : C \times C^\infty \rightarrow \mathbb{R}$, nonconstant, continuous, and linear, such that \succeq restricted to $(C \times C^\infty)^S$ is represented by $w(\cdot)$, where

$$w(F_{T-1}) = \int_{S_T} v(F_{T-1}(s_T)) dm_{T-1}, \quad F_{T-1} \in (C \times C^\infty)^S.$$

Since $u_2(\cdot)$ is continuous, linear and (by (B.4)) ordinally equivalent to $w(\cdot)$, it follows that $u_2(\cdot) = aw(\cdot) + b$ for some $a > 0$. By Risk Preference, it is without loss of generality to set $v(\ell) = \sum_0^\infty \delta^\tau u(\ell_\tau) = U'(\ell)$ for each $\ell \in C \times C^\infty$. Thus,

$$\mathcal{W}_{T-1}(c, F_{T-1}) = u_1(c) + a \int_{S_T} U'(F_{T-1}(s_T)) dm_{T-1} + b, \quad F_{T-1} \in (C \times C^\infty)^S.$$

Again by Risk Preference, it is without loss of generality to set $u_1(\cdot) = u(\cdot)$, $a = \delta$, and $b = 0$. State Independence, Risk Preference, and the nonconstancy of $u(\cdot)$ imply that m_{T-1} has full support.

To complete the proof of (B.2), extend the representation $\mathcal{W}_{T-1}(\cdot)$ of \succeq_{T-1} on $C_{T-1} \times (C \times C^\infty)^S$ to $C_{T-1} \times D_1$. We show that for every s_T , the preference $\succeq_{T-1|s_T}$ on $\mathcal{M}(C \times C^\infty)$ is *strategically rational*, that is, for any $M, N \in \mathcal{M}(C \times C^\infty)$, $M \succeq_{T-1|s_T} N$ implies $M \sim_{T-1|s_T} M \cup N$. Given Order, Continuity, Independence, and Set-Betweenness, the preference $\succeq_{T-1|s_T}$ has a (U, V) representation by **Kopylov (2007)**; given Risk Preference, State Independence, and Sophistication, $\succeq_{T-1|s_T}$ is non-trivial in that there exists $\{\ell, \ell'\} \in \mathcal{M}(C \times C^\infty)$ such that $\{\ell, \ell'\} \succ_{T-1|s_T} \{\ell'\}$. By Risk Preference, the restriction of

\succeq_T to $C \times C^\infty$ is represented by the function $\ell \mapsto \sum_0^\infty \delta^\tau u(\ell_\tau)$ and thus satisfies the vNM axioms. So by Sophistication and Lemma B.2, $U + V$ is ordinally equivalent to $\ell \mapsto \sum_0^\infty \delta^\tau u(\ell_\tau)$. But by State Independence and Risk Preference, U is ordinally equivalent to $\ell \mapsto \sum_0^\infty \delta^\tau u(\ell_\tau)$. Thus, V must be constant or ordinally equivalent to U . In either case, $\succeq_{T-1|s_T}$ must be strategically rational. Hence, for any $M \in \mathcal{M}(C \times C^\infty)$, there exists $\ell \in M$ that is \succeq_T -maximal in M , and for any c, s_T and $H \in D_1$,

$$(c, [H_{-s_T}, M]) \sim_{T-1} (c, [H_{-s_T}, \{\ell\}]).$$

This allows us to extend the representation $\mathcal{W}_{T-1}(\cdot)$ of \succeq_{T-1} on $C_{T-1} \times (C \times C^\infty)^S$ to $C_{T-1} \times D_1$ in the appropriate manner, completing the proof of (B.2).

As the induction hypothesis, suppose that for some $t < T$ and every τ satisfying $t \leq \tau < T - 1$, the restriction of $\succeq_{\tau+1}$ to $C_{\tau+1} \times D_{T-\tau-1}$ is represented by

$$\mathcal{W}_{\tau+1}(c, F_{\tau+1}) = u(c) + \delta \int_{S_{\tau+2}} U_{\tau+2}(F_{\tau+1}(s_{\tau+2}), s_{\tau+2}) dm_{\tau+1} \quad (c, F_{\tau+1}) \in C_{\tau+1} \times D_{T-\tau-1},$$

where $m_{\tau+1}$ has full support, $U_{\tau+2}(\cdot, s_{\tau+2}) : \mathcal{M}(C_{\tau+2} \times D_{T-\tau-2}) \rightarrow \mathbb{R}^1$ is nonconstant, continuous, and linear and is defined recursively via

$$\begin{aligned} &U_{\tau+2}(M_{\tau+2}, s_{\tau+2}) \\ &= \max_{(c, F_{\tau+2}) \in M_{\tau+2}} \left\{ u(c) + \delta \int_{S_{\tau+3}} U_{\tau+3}(F_{\tau+2}(s_{\tau+3}), s_{\tau+3}) dp_{\tau+2} \right. \\ &\quad \left. + \frac{1 - \alpha_{\tau+2}}{\alpha_{\tau+2}} \left(u(c) + \delta \int_{S_{\tau+3}} U_{\tau+3}(F_{\tau+2}(s_{\tau+3}), s_{\tau+3}) dq_{\tau+2} \right) \right\} \\ &\quad - \max_{(c', F'_{\tau+2}) \in M_{\tau+2}} \frac{1 - \alpha_{\tau+2}}{\alpha_{\tau+2}} \left\{ u(c') + \delta \int_{S_{\tau+3}} U_{\tau+3}(F'_{\tau+2}(s_{\tau+3}), s_{\tau+3}) dq_{\tau+2} \right\}, \end{aligned}$$

and the boundary condition

$$U_{T-1}(M_{T-1}, s_T) = U^T(M_{T-1}), \quad M_{T-1} \in \mathcal{M}(C \times C^\infty).$$

Moreover,

$$\begin{aligned} &\alpha_{\tau+2} \in (0, 1], p_{\tau+2}, q_{\tau+2} \in \Delta(S_{\tau+2}), \text{ each } p_{\tau+2} \text{ has full support,} \\ &\text{and } m_{\tau+2} = \alpha_{\tau+2} p_{\tau+2} + (1 - \alpha_{\tau+2}) q_{\tau+2}. \end{aligned}$$

We construct \mathcal{W}_t having the appropriate form and representing \succeq_t .²¹ The argument is divided into a series of steps.

²¹For $t = 0$, the measure m_0 over S_1 that we construct can be denoted instead by p_0 , as in the desired representation.

STEP 1. Define the “convex hull” of contingent menus.

For any mixture space, we have the usual notion of convex hull of a set M —the smallest convex (mixture-closed) set containing M . However, a mixture space framework is not adequate because, for example, $\mathcal{M}(C_T \times \mathcal{C}_T)$ is not a mixture space: $\lambda[\lambda'M + (1 - \lambda')M'] + (1 - \lambda)M' \neq \lambda\lambda'M + (1 - \lambda\lambda')M'$ if M and M' are not convex. More generally, because $\alpha M + (1 - \alpha)M \neq M$ in general, the “convex hull” of any M need not contain M . In fact, we are interested in the convex hull of contingent menus. Thus we define $\text{co}(F_t)$ for any F_t in \mathcal{C}_t and we do so by backward induction.

Since $C_T \times C^\infty$ is a mixture space, the “convex hull of $M_{T-1} \in \mathcal{M}(C_T \times C^\infty)$ ” has the usual meaning—the smallest convex set containing M_{T-1} . For any contingent menu F_{T-1} in D_1 , define its convex hull $\text{co}(F_{T-1})$ as the contingent menu that maps each s_T into $\text{co}(F_{T-1}(s_T))$. Let

$$\mathcal{D}_1 = \{\text{co}(F'_{T-1}) : F'_{T-1} \in D_1\} \subset D_1.$$

Then \mathcal{D}_1 is a mixture space.

For the inductive step, suppose that $\text{co}(\cdot)$ has been defined on D_{T-t-1} , and that

$$\mathcal{D}_{T-t-1} = \{\text{co}(F'_{t+1}) : F'_{t+1} \in D_{T-t-1}\} \subset D_{T-t-1}$$

is a mixture space. Let $F_t \in D_{T-t}$, $s_{t+1} \in S_{t+1}$, and

$$N = \{(c_{t+1}, \text{co}(F_{t+1})) : (c_{t+1}, F_{t+1}) \in F_t(s_{t+1})\}.$$

Since $C_{t+1} \times \mathcal{D}_{T-t-1}$ is a mixture space, the smallest convex subset of $C_{t+1} \times \mathcal{D}_{T-t-1}$ containing N is well-defined. We define $\text{co}(F_t)(s_{t+1})$ to be that set. This defines $\text{co}(F_t)$. Note that it lies in $\mathcal{D}_{T-t} = \{\text{co}(F'_t) : F'_t \in D_{T-t}\}$, and that the latter is a mixture space.

STEP 2. Each \succeq_t satisfies Indifference to Randomization, that is,

$$(c, F_t) \sim_t (c, \text{co}(F_t)). \tag{B.5}$$

PROOF. For $t = T - 1$, since $C \times C^\infty$ is a mixture space, the preference \succeq_{T-1} restricted to $C_{T-1} \times D_1 = C_{T-1} \times (\mathcal{M}(C \times C^\infty))^{S_T}$ satisfies IR by Order, Continuity, and Independence (see Dekel et al. 2001, Lemma 1).²²

However, $C_t \times D_{T-t}$ is not a mixture space if $t < T - 1$. Fortunately, we can invoke Kopylov (2007) to prove (B.5).²³ He extends the GP theorem to a domain, consisting of hierarchies of menus, that corresponds to our setting when the state space S is a singleton and when consumption occurs only at the terminal time. His arguments are readily adapted to accommodate the multiplicity of states and the presence of intermediate consumption. ◁

²²Their result is formulated for preference defined on menus of lotteries, but the same argument can be used for menus of any compact metric mixture space. The contingent nature of menus in our case is of no significance because mixtures are defined statewise.

²³We are grateful to Igor Kopylov for pointing out this line of argument.

STEP 3. The order \succeq_t restricted to $C_t \times D_{T-t}$ can be represented by $\widehat{\mathcal{W}}_t(\cdot)$ having the form

$$\widehat{\mathcal{W}}_t(c, F) = u_t^*(c) + \sum_{s_{t+1}} U_{t+1}^*(F(s_{t+1}), s_{t+1}), \tag{B.6}$$

where $u_t^*(\cdot)$ and $U_{t+1}^*(\cdot, s_{t+1})$ are nonconstant, continuous, and linear on C_t and $\mathcal{M}(C_{t+1} \times D_{T-t-1})$ respectively, and where

$$U_{t+1}^*(M, s_{t+1}) = U_{t+1}^*(\text{co}(M), s_{t+1}) \quad \text{for } M \in \mathcal{M}(C_{t+1} \times D_{T-t-1}). \tag{B.7}$$

PROOF. To prove this, restrict attention first to $C_t \times \mathcal{D}_{T-t}$. Each F in \mathcal{D}_{T-t} maps S_{t+1} into $\mathcal{M}^c(C_{t+1} \times \mathcal{D}_{T-t-1})$, the collection of convex (and closed) subsets of the mixture space $C_{t+1} \times \mathcal{D}_{T-t-1}$. But $\mathcal{M}^c(C_{t+1} \times \mathcal{D}_{T-t-1})$ is a mixture space. Since \succeq_t satisfies Order, Continuity, and Independence on $C_t \times \mathcal{D}_{T-t}$, it admits a utility representation there by some $\widehat{\mathcal{W}}_t : C_t \times \mathcal{D}_{T-t} \rightarrow \mathbb{R}^1$ having the form (B.6); additivity across c and F can be established as in (B.3), while the additive separability across states follows as in **Kreps (1988, Propn. 7.4)**, for example. Use (B.7) to extend (B.6) to all of $C_t \times D_{T-t}$. Indifference to Randomization (Step 2) implies that $\widehat{\mathcal{W}}_t(\cdot)$ represents \succeq_t on $C_t \times D_{T-t}$. \triangleleft

Let $\succeq_{t|s_{t+1}}$ on $\mathcal{M}(C_{t+1} \times D_{T-t-1})$ be the preference represented by $U_{t+1}^*(\cdot, s_{t+1})$.

STEP 4. $\succeq_{t|s_{t+1}}$ satisfies the GP axioms suitably translated to $\mathcal{M}(C_{t+1} \times D_{T-t-1})$.

Thus by GP's theorem and the extension provided by **Kopylov (2007)**,²⁴

$$U_{t+1}^*(M, s_{t+1}) = \max_{(c, F) \in M} \{U_{t+1}^{GP}(c, F, s_{t+1}) + V_{t+1}^{GP}(c, F, s_{t+1})\} - \max_{(c', F') \in M} V_{t+1}^{GP}(c', F', s_{t+1}),$$

for some $U_{t+1}^{GP}(\cdot, s_{t+1})$ and $V_{t+1}^{GP}(\cdot, s_{t+1})$, continuous and linear functions on $C_{t+1} \times D_{T-t-1}$. The subscript t indicates that these functions may depend also on the history s_t^1 underlying \succeq_t .

STEP 5. For some $A(s_{t+1}) > 0$,

$$\begin{aligned} &U_{t+1}^{GP}(c, F, s_{t+1}) + V_{t+1}^{GP}(c, F, s_{t+1}) \\ &= A(s_{t+1}) \left[u(c) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dm_{t+1} \right]. \end{aligned} \tag{B.8}$$

PROOF. By Risk Preference and State Independence, for any c, H, s_{t+1} there exist $\ell, \ell' \in \mathcal{L}_{t+1}$ such that

$$(c, [H_{-s_{t+1}}, \{\ell\}]) \succ_t (c, [H_{-s_{t+1}}, \{\ell'\}]) \text{ and } \ell \succ_{t+1} \ell'.$$

²⁴GP work with a domain of menus of lotteries. Their theorem would apply directly if we had adopted the larger domain obtained by replacing (4) with $F_t : S_{t+1} \rightarrow \mathcal{M}(\Delta(C_{t+1} \times D_{T-t-1}))$. However, adding an extra layer of lotteries can be avoided by invoking Kopylov, suitably extended to accommodate a finite (nonsingleton) state space and intermediate consumption. (His Temporal Set-Betweenness axiom is satisfied by our preference $\succeq_{t|s_{t+1}}$, by **Lemma B.4** and Set-Betweenness.)

It follows from Sophistication that

$$(c, [H_{-s_{t+1}}, \{\ell, \ell'\}]) \succ_t (c, [H_{-s_{t+1}}, \{\ell'\}]).$$

In particular, the preference $\succeq_t|_{s_{t+1}}$ on $\mathcal{M}(C_{t+1} \times D_{T-t-1})$ satisfies $\{\ell, \ell'\} \succ_t|_{s_{t+1}} \{\ell'\}$. By **Step 4** this preference has a $(U_{t+1}^{GP}, V_{t+1}^{GP})$ representation, and thus by Sophistication, Order, Continuity, and Independence for \succeq_{t+1} , **Lemma B.2** implies that $U_{t+1}^{GP}(\cdot, s_{t+1}) + V_{t+1}^{GP}(\cdot, s_{t+1})$ represents \succeq_{t+1} . By the induction hypothesis, \succeq_{t+1} is represented also by $\mathcal{W}_{t+1}(\cdot)$, and since both functions are continuous and linear, they must be cardinally equivalent. Thus **(B.8)** follows. \triangleleft

STEP 6. Let $V_{t+1}(c, F, s_{t+1}) = (1/A(s_{t+1}))V_{t+1}^{GP}(c, F, s_{t+1})$. Then

$$V_{t+1}(c, F, s_{t+1}) = w_{t+1}(c, s_{t+1}) + \sum_{s_{t+2}} v_{t+1}(F(s_{t+2}), s_{t+1}, s_{t+2}), \tag{B.9}$$

where $w_{t+1}(\cdot, s_{t+1})$ and each $v_{t+1}(\cdot, s_{t+1}, s_{t+2})$ are continuous and linear on C_{t+1} and $\mathcal{M}(C_{t+2} \times D_{T-t-2})$ respectively.

PROOF. The function $M \mapsto V_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1})$ gives the (temptation) utility of the indicated consumption and contingent menu pair as a function of the menu M provided in state s_{t+2} . Similarly for the function $M \mapsto \bar{U}_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1})$, where

$$\bar{U}_{t+1}(c, F, s_{t+1}) = \frac{1}{A(s_{t+1})} U_{t+1}^{GP}(c, F, s_{t+1}).$$

Recall the order $\succeq_t|_{s_{t+1}, s_{t+2}}$ defined prior to **Lemma B.3**. For any given c and F , it is represented by

$$L \mapsto \max_{M \in L} \{ \bar{U}_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1}) + V_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1}) \} \\ - \max_{M' \in L} V_{t+1}(c, [F_{-s_{t+2}}, M'], s_{t+1}),$$

for any closed $L \subset \mathcal{M}(C_{t+2} \times D_{T-t-2})$. By Risk Preference, State Independence, and **Lemma B.4**, $\bar{U}_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1})$ is nonconstant, and so by **Lemma B.3**, $\succeq_t|_{s_{t+1}, s_{t+2}}$ is strategically rational. By **Lemma B.1(i)**, if $V_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1})$ is nonconstant then it is ordinally equivalent to $\bar{U}_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1}) + V_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1})$, which by **Step 5** is ordinally equivalent to $U_{t+2}(\cdot, s_{t+2})$. Thus, if $V_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1})$ is nonconstant, then for all $M, M' \in \mathcal{M}(C_{t+2} \times D_{T-t-2})$,

$$V_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1}) \geq V_{t+1}(c, [F_{-s_{t+2}}, M'], s_{t+1}) \tag{B.10} \\ \iff U_{t+2}(M, s_{t+2}) \geq U_{t+2}(M', s_{t+2}) \\ \iff U_{t+2}(\text{co}(M), s_{t+2}) \geq U_{t+2}(\text{co}(M'), s_{t+2}) \\ \iff V_{t+1}(c, [F_{-s_{t+2}}, \text{co}(M)], s_{t+1}) \geq V_{t+1}(c, [F_{-s_{t+2}}, \text{co}(M')], s_{t+1}),$$

where use has been made of (B.7). On the other hand, if $V_{t+1}(c, [F_{-s_{t+2}}, \cdot], s_{t+1})$ is constant, then the equivalence of the first and last lines is clear. We conclude that for every F , c , and s_{t+2} ,

$$V_{t+1}(c, [F_{-s_{t+2}}, M], s_{t+1}) = V_{t+1}(c, [F_{-s_{t+2}}, \text{co}(M)], s_{t+1}).$$

Repeated application of this equality for all states in S_{t+2} yields

$$V_{t+1}(c, F, s_{t+1}) = V_{t+1}(c, \text{co}(F), s_{t+1}),$$

a form of indifference to randomization for V_{t+1} . Thus one can argue as in Step 3 to derive (B.9). \triangleleft

STEP 7. For some $\gamma(s_{t+1}) \geq 0$, continuous linear function $w(\cdot, s_{t+1})$ on C_{t+1} and $q_{t+1} \in \Delta(S_{t+2})$,

$$V_{t+1}^{GP}(c, F, s_{t+1}) = A(s_{t+1}) \left[w_{t+1}(c, s_{t+1}) + \gamma(s_{t+1}) \int_{S_{t+2}} U_{t+2}(M, s_{t+2}) dq_{t+1}(s_{t+2}) \right].$$

PROOF. We begin by providing structure on each $v_{t+1}(\cdot, s_{t+1}, s_{t+2})$ in (B.9): we show that

$$v_{t+1}(\cdot, s_{t+1}, s_{t+2}) = a(s_{t+1}, s_{t+2})U_{t+2}(\cdot, s_{t+2}) + b(s_{t+1}, s_{t+2}), \tag{B.11}$$

for some $a(s_{t+1}, s_{t+2}) \geq 0$. Given (B.9), we can refine (B.10) into the statement that if $v_{t+1}(\cdot, s_{t+1}, s_{t+2})$ is nonconstant, then $v_{t+1}(\cdot, s_{t+1}, s_{t+2})$ is ordinally equivalent to $U_{t+2}(\cdot, s_{t+2})$. Given continuity and linearity of both functions, (B.11) holds for some $a(s_{t+1}, s_{t+2}) > 0$. If $v_{t+1}(\cdot, s_{t+1}, s_{t+2})$ is constant, then (B.11) holds with $a(s_{t+1}, s_{t+2}) = 0$.

Define $\gamma(s_{t+1})$ and the measure q_{t+1} over S_{t+2} by

$$\begin{aligned} \gamma(s_{t+1}) &= \sum_{S_{t+2}} a(s_{t+1}, s_{t+2}) \geq 0, \\ q_{t+1}(s_{t+2}) &= \begin{cases} \frac{a(s_{t+1}, s_{t+2})}{\gamma(s_{t+1})} & \text{if } \gamma(s_{t+1}) > 0 \\ m_{t+1}(s_{t+2}) & \text{otherwise.} \end{cases} \end{aligned}$$

Then

$$V_{t+1}(c, F, s_{t+1}) = w_{t+1}(c, s_{t+1}) + \gamma(s_{t+1}) \int_{S_{t+2}} U_{t+2}(M, s_{t+2}) dq_{t+1}(s_{t+2}) + k,$$

where $k = \sum_{S_{t+2}} b(s_{t+1}, s_{t+2})$. Set $k = 0$ without loss of generality. \triangleleft

STEP 8. For some $0 < \alpha_{t+1}(s_{t+1}) \leq 1$,

$$V_{t+1}^{GP}(c, F, s_{t+1}) = A(s_{t+1})(1 - \alpha_{t+1}(s_{t+1})) \left(u(c) + \delta \int_{S_{t+2}} U_{t+2}(F(s_{t+2}), s_{t+2}) dq_{t+1} \right). \tag{B.12}$$

PROOF. By Risk Preference and State Independence, $U_{t+1}^{GP}(\ell, s_{t+1})$ is ordinally (and hence cardinally) equivalent to the continuous linear function $\ell \mapsto \sum_{\tau=t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau)$. Thus without loss of generality

$$U_{t+1}^{GP}(\ell, s_{t+1}) = A(s_{t+1})\alpha_{t+1}(s_{t+1}) \left[\sum_{\tau=t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau) \right]$$

for some $\alpha_{t+1}(s_{t+1}) > 0$. By **Step 5**,

$$U_{t+1}^{GP}(\ell, s_{t+1}) + V_{t+1}^{GP}(\ell, s_{t+1}) = A(s_{t+1}) \left[\sum_{\tau=t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau) \right].$$

Thus, $V_{t+1}^{GP}(\ell, s_{t+1}) = A(s_{t+1})(1 - \alpha_{t+1}(s_{t+1})) \sum_{\tau=t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau)$, and by **Step 7**,

$$\begin{aligned} w_{t+1}(\ell_{t+1}, s_{t+1}) + \gamma(s_{t+1}) \sum_{\tau=t+2}^{T+1} \delta^{\tau-(t+2)} u(\ell_\tau) \\ = (1 - \alpha_{t+1}(s_{t+1})) \left(u(\ell_{t+1}) + \delta \sum_{\tau=t+2}^{T+1} \delta^{\tau-(t+2)} u(\ell_\tau) \right) \end{aligned}$$

$$\implies w_{t+1}(\ell_{t+1}, s_{t+1}) - (1 - \alpha_{t+1}(s_{t+1}))u(\ell_{t+1})$$

$$= [(1 - \alpha_{t+1}(s_{t+1}))\delta - \gamma(s_{t+1})] \sum_{\tau=t+2}^{T+1} \delta^{\tau-(t+2)} u(\ell_\tau).$$

Since $u(\cdot)$ is nonconstant, we have $(1 - \alpha_{t+1}(s_{t+1}))\delta = \gamma(s_{t+1})$ and $w_{t+1}(\ell_{t+1}, s_{t+1}) = (1 - \alpha_{t+1}(s_{t+1}))u(\ell_{t+1})$. If $\gamma(s_{t+1}) = 0$, then $\delta > 0$ implies $w_{t+1}(\ell_{t+1}, s_{t+1}) = 0$, which yields (B.12) with $\alpha_{t+1}(s_{t+1}) = 1$. On the other hand, if $\gamma(s_{t+1}) > 0$, then $\delta > 0$ implies (B.12) with $\alpha_{t+1}(s_{t+1}) < 1$. \triangleleft

STEP 9. The unique measure p_{t+1} over S_{t+2} satisfying $m_{t+1} = \alpha_{t+1}p_{t+1} + (1 - \alpha_{t+1})q_{t+1}$ is a probability measure with full support and furthermore

$$U_{t+1}^{GP}(c, F, s_{t+1}) = A(s_{t+1})\alpha_{t+1}(s_{t+1}) \left(u(c) + \delta \int_{S_{t+2}} U_{t+2}(F(s_{t+2}), s_{t+2}) dp_{t+1} \right). \quad (\text{B.13})$$

PROOF. Steps 5 and 8 yield (B.13), given that p_{t+1} satisfies $m_{t+1} = \alpha_{t+1}p_{t+1} + (1 - \alpha_{t+1})q_{t+1}$. We show next that p_{t+1} is a probability measure with full support. The definition of p_{t+1} implies that $\sum_{s_{t+2}} p_{t+1}(s_{t+2}) = 1$. To see that $p_{t+1}(s_{t+2}) > 0$ for all s_{t+2} , note that $U_{t+2}(\cdot, s_{t+2})$ is nonconstant (by the induction hypothesis) and that for any $s_{t+1}, s_{t+2}, c', c, M', M, F$, and G ,

$$\begin{aligned} U_{t+2}(M', s_{t+2}) \geq U_{t+2}(M, s_{t+2}) &\iff (c, [F_{-s_{t+2}}, M']) \succeq_{t+1} (c, [F_{-s_{t+2}}, M]) \\ &\iff^* (c', [G_{-s_{t+1}}, \{(c, [F_{-s_{t+2}}, M'])\}]) \succeq_t (c', [G_{-s_{t+1}}, \{(c, [F_{-s_{t+2}}, M])\}]) \\ &\iff U_{t+1}^{GP}(c, [F_{-s_{t+2}}, M'], s_{t+1}) \geq U_{t+1}^{GP}(c, [F_{-s_{t+2}}, M], s_{t+1}) \\ &\iff U_{t+2}(M', s_{t+2})p_{t+1}(s_{t+2}) \geq U_{t+2}(M, s_{t+2})p_{t+1}(s_{t+2}), \end{aligned}$$

where the equivalence \iff^* is implied by **Lemma B.4**. \triangleleft

STEP 10. *Complete the inductive step.*

Since $A_t(s_{t+1})\alpha_{t+1}(s_{t+1}) > 0$ for all s_{t+1} , we have $\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1}) > 0$. Consider the positive affine transformation of $\widehat{\mathcal{W}}_t$ defined by

$$\begin{aligned} \mathcal{W}_t(c, F) &= \frac{\delta}{\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} \widehat{\mathcal{W}}_t(c, F) \\ &= \frac{\delta}{\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} u_{t+1}^*(c) \\ &\quad + \frac{\delta}{\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} \sum_{s_{t+1}} U_{t+1}^*(F(s_{t+1}), s_{t+1}), \end{aligned}$$

for all $(c, F) \in C_t \times D_{T-t}$. Obviously, $\mathcal{W}_t(\cdot)$ represents \succeq_t on $C_t \times D_{T-t}$. Define

$$\begin{aligned} u_{t+1}(c) &\equiv \frac{\delta}{\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} u_{t+1}^*(c), \\ U_{t+1}(M_{t+1}, s_{t+1}) &\equiv \frac{1}{A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} U_{t+1}^*(M, s_{t+1}), \\ m_t(s_{t+1}) &= \frac{A_t(s_{t+1})\alpha_{t+1}(s_{t+1})}{\sum_{s_{t+1}} A_t(s_{t+1})\alpha_{t+1}(s_{t+1})} > 0 \text{ for each } s_{t+1}. \end{aligned}$$

Then m_t has full support and

$$\mathcal{W}_t(c, F) = u_{t+1}(c) + \delta \int_{S_{t+1}} U_{t+1}(F(s_{t+1}), s_{t+1}) dm_t(s_{t+1}), \quad F_t \in D_{T-t},$$

where

$$\begin{aligned} &U_{t+1}(M_{t+1}, s_{t+1}) \\ &= \max_{(c, F_{t+1}) \in M_{t+1}} \left\{ u(c) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dp_{t+1} \right. \\ &\quad \left. + \frac{(1 - \alpha_{t+1})}{\alpha_{t+1}} \left(u(c) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dq_{t+1} \right) \right\} \\ &\quad - \max_{(c', F'_{t+1}) \in M_{t+1}} \left\{ \frac{(1 - \alpha_{t+1})}{\alpha_{t+1}} \left(u(c') + \delta \int_{S_{t+2}} U_{t+2}(F'_{t+1}(s_{t+2}), s_{t+2}) dq_{t+1}(s_{t+2}) \right) \right\}. \end{aligned}$$

It remains to show that $u_{t+1}(\cdot) = u(\cdot)$. By Risk Preference and the representation $\mathcal{W}_t(\cdot)$, the following functions are ordinally equivalent:

$$\begin{aligned} (c, \ell) &\longmapsto u(c) + \delta \left[\sum_{t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau) \right] \\ (c, \ell) &\longmapsto u_{t+1}(c) + \delta \left[\sum_{t+1}^{T+1} \delta^{\tau-(t+1)} u(\ell_\tau) \right]. \end{aligned}$$

Since both are continuous linear functions, they must be cardinally equivalent. An argument analogous to that used in **Step 8** yields the desired result, and completes the inductive step.

We now have a representation for \succeq_0 on $C \times D_T$ for each T . It remains to extend the representation to $C \times \mathcal{C}$. Note that $D_T \subset D_{T+1}$ for each T . Thus each representation on $C \times D_{T+1}$ induces a representation also on $C \times D_T$. We proceed by showing: (i) consistency of the representations of \succeq_0 obtained above, (ii) a (unique) continuous extension to $C \times \mathcal{C}$, and (iii) the extension has the appropriate functional form (5)–(7). The last step simultaneously derives the desired representation of (\succeq_t) via **Lemma B.2**.

To show (i), consider the representations $(p_0, (p_t, q_t, \alpha_t)_1^T)$ and $(p'_0, (p'_t, q'_t, \alpha'_t)_1^{T+1})$ of \succeq_0 on $C \times D_T$ and $C \times D_{T+1}$, respectively. We need to show that $p_0 = p'_0$ and for all $0 < t \leq T$, $p_t = p'_t$, $q_t = q'_t$, and $\alpha_t = \alpha'_t$. This can be proved by adapting the argument used in **Epstein (2006, Corollary 3.3)**. The uniqueness part of the Anscombe–Aumann theorem ensures $p_0 = p'_0$. Further, if \succeq_0 exhibits a preference for commitment conditional on a history of length $t - 1 \leq T$, then it must be that $p_t = p'_t$, $q_t = q'_t$, and $\alpha_t = \alpha'_t$. If there is no such preference for commitment, then the non-uniqueness of the representation permits us to set $p_t = p'_t$, $q_t = q'_t$, and $\alpha_t = \alpha'_t$ without loss of generality. Proceeding in this way, we obtain $(p_0, (p_t, q_t, \alpha_t)_1^\infty)$ such that \succeq_0 has the desired representation $\mathcal{W}_0 : C \times \cup_1^\infty D_T \rightarrow \mathbb{R}$.

For (ii), we exploit the denseness indicated in **Theorem A.1(iv)**. First, observe that $\mathcal{W}_0(\cdot)$ is bounded above by $\mathcal{W}_0(\bar{\ell}) = (1 - \delta)^{-1} \max_{c \in C} u(c)$ and below by $\mathcal{W}_0(\underline{\ell}) = (1 - \delta)^{-1} \min_{c \in C} u(c)$, and in particular, each $(c, F) \in C \times \mathcal{C}$ is ranked between the risky streams $\bar{\ell}, \underline{\ell} \in \mathcal{L}$. By Continuity and Risk Preference it follows that each $(c, F) \in C \times \mathcal{C}$ is indifferent to some unique mixture of $\bar{\ell}, \underline{\ell}$, which we denote by $\bar{\ell} \lambda_{(c,F)} \underline{\ell} \in \mathcal{L}$. Since $\mathcal{L} \subset C \times \cup_1^\infty D_T$, this allows us to define an extension of $\mathcal{W}_0(\cdot)$ to all of $C \times \mathcal{C}$ by setting $\mathcal{W}_0(c, F) = \mathcal{W}_0(\bar{\ell} \lambda_{(c,F)} \underline{\ell})$. To see that this extension is continuous, without loss of generality let $\mathcal{W}_0(\bar{\ell}) = 1$ and $\mathcal{W}_0(\underline{\ell}) = 0$ so that in fact $\mathcal{W}_0(c, F) = \lambda_{(c,F)}$ and suppose $(c_n, F_n) \rightarrow (c, F)$ and, by way of contradiction, $\lambda_{(c_n, F_n)} \not\rightarrow \lambda_{(c,F)}$. Then for some ε -ball $B(\lambda_{(c,F)}, \varepsilon)$ around $\lambda_{(c,F)}$, there are infinitely many n such that $\lambda_{(c_n, F_n)} \notin B(\lambda_{(c,F)}, \varepsilon)$. Let (c_m, F_m) denote the corresponding subsequence and note that $(c_m, F_m) \rightarrow (c, F)$. Since $\{\lambda_{(c_m, F_m)}\}$ is a subsequence in the unit interval, it has a convergent subsequence $\{\lambda_{(c_i, F_i)}\}$ with a limit different from $\lambda_{(c,F)}$; denote by (c_i, F_i) the corresponding subsequence of (c_m, F_m) . Then Continuity implies that $(c, F) = \lim(c_i, F_i) \sim \lim \bar{\ell} \lambda_{(c_i, F_i)} \underline{\ell} \neq \bar{\ell} \lambda_{(c,F)} \underline{\ell}$. However, $(c, F) \sim \bar{\ell} \lambda_{(c,F)} \underline{\ell}$ for a unique $\lambda_{(c,F)}$, a contradiction. Thus we have a continuous extension of $\mathcal{W}_0(\cdot)$ to $C \times \mathcal{C}$. Since the latter is compact, the extension is uniformly continuous and unique.

For (iii), uniform continuity of \mathcal{W}_0 is the key. It implies uniform continuity of $U_1(\cdot, s_1)$ on $\mathcal{M}(C \times \cup_1^\infty D_T)$, and also of $\mathcal{Q}_1(\cdot)$ and $\mathcal{V}_1(\cdot)$ on $C \times \cup_1^\infty D_T$; as a result the latter two functions can be extended uniquely to continuous functions on $C \times \mathcal{C}$. Argue inductively. The details are tedious but straightforward.

This completes the proof of sufficiency. The proof for uniqueness is similar to that in **Epstein (2006)**, and thus is omitted.

REFERENCES

- Aliprantis, Charalambos D. and Kim C. Border (1994), *Infinite-Dimensional Analysis*. Springer, Berlin. [196]
- Brandenburger, Adam and Eddie Dekel (1993), “Hierarchies of beliefs and common knowledge.” *Journal of Economic Theory*, 59, 189–198. [214]
- Camerer, Colin (1995), “Individual decision making.” In *The Handbook of Experimental Economics* (John H. Kagel and Alvin E. Roth, eds.), 587–703, Princeton University Press, Princeton. [195]
- Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini (2001), “Representing preferences with a unique subjective state space.” *Econometrica*, 69, 891–934. [221]
- Epstein, Larry G. (1983), “Stationary cardinal utility and optimal growth under uncertainty.” *Journal of Economic Theory*, 31, 133–152. [210]
- Epstein, Larry G. (2006), “An axiomatic model of non-Bayesian updating.” *Review of Economic Studies*, 73, 413–436. [193, 198, 201, 207, 208, 212, 227]
- Epstein, Larry G., Jawwad Noor, and Alvaro Sandroni (2008), “Supplementary appendix for ‘Non-Bayesian updating: A theoretical framework.’” Working Paper 08-017, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania. [195]
- Epstein, Larry G. and Michael Peters (1999), “A revelation principle for competing mechanisms.” *Journal of Economic Theory*, 88, 119–160. [214]
- Epstein, Larry G. and Tan Wang (1996), “‘Beliefs about beliefs’ without probabilities.” *Econometrica*, 64, 1343–1373. [214, 216]
- Gul, Faruk and Wolfgang Pesendorfer (2001), “Temptation and self-control.” *Econometrica*, 69, 1403–1435. [194, 217]
- Gul, Faruk and Wolfgang Pesendorfer (2004), “Self-control and the theory of consumption.” *Econometrica*, 72, 119–158. [194, 214]
- Kopylov, Igor (2007), “Finite additive utility representations for preferences over menus.” Unpublished paper, Department of Economics, University of California, Irvine. [194, 219, 221, 222]
- Kreps, David M. (1988), *Notes on the Theory of Choice*. Westview, Boulder. [198, 208, 222]
- Mertens, Jean-François and Shmuel Zamir (1985), “Formulation of Bayesian analysis for games with incomplete information.” *International Journal of Game Theory*, 14, 1–29. [214]

Mullainathan, Sendhil (2000), "Thinking through categories." Unpublished paper, Department of Economics, MIT. [195]

Noor, Jawwad (2008), "Temptation, welfare and revealed preference." Unpublished paper, Department of Economics, Boston University. [201, 217]

Rabin, Matthew (1998), "Psychology and economics." *Journal of Economic Literature*, 36, 11–46. [195]

Rabin, Matthew (2002), "Inference by believers in the law of small numbers." *Quarterly Journal of Economics*, 117, 775–816. [195, 204]

Strotz, Robert H. (1955), "Myopia and inconsistency in dynamic utility maximization." *Review of Economic Studies*, 23, 165–180. [201]

Tversky, Amos and Daniel Kahneman (1974), "Judgement under uncertainty: Heuristics and biases." *Science*, 185, 1124–1131. [195]

Submitted 2006-6-9. Final version accepted 2008-4-9. Available online 2008-4-9.