

Rochester Center for

Economic Research

Non-Bayesian Updating : A Theoretical Framework

Larry G. Epstein and Alvaro Sandroni

Working Paper No. 505
December 2003

UNIVERSITY OF
ROCHESTER

NON-BAYESIAN UPDATING: A THEORETICAL FRAMEWORK*

Larry G. Epstein Alvaro Sandroni

April 2004

Abstract

This paper models an agent in an infinite horizon 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. Then the model is specialized axiomatically to capture updating biases that reflect excessive weight given to (i) prior beliefs, or alternatively, (ii) the realized sample. Finally, the paper describes a counterpart of the exchangeable Bayesian model, where the agent tries to learn about parameters, and some answers are provided to the question "what does a non-Bayesian updater learn?"

1. INTRODUCTION

This paper models an agent in an infinite horizon setting who does not update according to Bayes' Rule, and who is self-aware and anticipates her updating behavior when formulating plans. Three central questions are addressed.

Are there axiomatic foundations for such a model? We provide such foundations in the form of a representation theorem for suitably defined preferences. A dynamic version of the (Savage or) Anscombe-Aumann theorem provides the foundation

*Department of Economics, University of Rochester, Rochester, NY 14627, lepn@troi.cc.rochester.edu and J. L. Kellogg School of Management, MEDS, Northwestern University, Evanston, IL 60208, sandroni@kellogg.northwestern.edu. We have benefitted from conversations with Joseph Perktold and Werner Ploberger.

for reliance on a probability measure representing prior beliefs and for subsequent Bayesian updating of the prior belief as information arrives. We generalize this Anscombe-Aumann theorem so that *both* the prior *and* the way in which it is updated are subjective, that is, are derived from preference. The model is dynamic: consumption processes are the ultimate source of utility, conditional preferences are defined at each time-history pair and these preferences are dynamically consistent. The latter implies that dynamic choice behavior is derived from preferences at time 0, as in the Bayesian model. Thus, though the model is not normative, the agent is rational in the sense of maximizing a stable, transitive and complete preference relation.¹

What updating rules are permitted? Our general framework is rich: 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. Richness is demonstrated by axiomatic specializations that capture excessive weight given at the updating stage to (i) prior beliefs, or alternatively, (ii) the realized sample. A counterpart of the exchangeable Bayesian model, where the agent tries to learn about parameters, is also described.

It will be evident that there are many other kinds of updating biases that can be accommodated, including biases similar to some that have been observed in experimental psychology; see Tversky and Kahneman [23] and the surveys by Camerer [3] and Rabin [18], for example. Our model does not address the experimental evidence directly, however, because the latter deals with the updating of objective probabilities, while our model makes sense only when probabilities are subjective.²

What do non-Bayesian updaters learn? A central focus of the literature on Bayesian learning is on what is learned asymptotically and how an agent forecasts as more and more observations are available. Bayesian forecasts are eventually correct with probability 1 under the truth given suitable conditions, the key condition being absolute continuity of the true measure with respect to initial beliefs.

¹As explained in Section 2.1, this claim must be qualified.

²See [19] and [17], for example, for models of updating for objective probabilities that address the experimental evidence. Though the associated models of preference are not made explicit, to the best of our understanding these authors assume implicitly that the agent is an expected utility maximizer who is naive in the sense of not anticipating future deviations from Bayesian updating nor the fact that today's plans may not be implemented. In contrast, our agent is sophisticated and dynamically consistent.

Hence, multiple repetitions of Bayes' Rule transforms the historical record into a near perfect guide for the future. We investigate the corresponding question for non-Bayesian updaters who face a statistical inference problem and conform to one of the above noted biases. We describe simple non-Bayesian updating rules that, if repeated multiple times, will also uncover the true data generating process. However, our richer hypothesis about updating behavior permits a broader range of possibilities for what is learned in the long run. In one of our results, we show that convergence to correct forecasts holds for an agent who underreacts to observations when updating. If she overreacts then her forecasts are eventually correct with positive probability - an example shows that with positive probability she may become certain that a false parameter is true and thus converge to precise but false forecasts.

The issue of foundations for non-Bayesian updating is taken up in [5] in a three-period framework, where the agent updates once and consumption occurs only at the terminal time. The model is extended here to an infinite horizon setting. We take as the benchmark the standard specification of utility in dynamic modeling, whereby utility at time t is given by

$$U_t(c) = E_t \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} u(c_{\tau}) \right], \quad t = 0, 1, \dots, \quad (1.1)$$

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 that is updated by Bayes' Rule. Our model generalizes (1.1) to which it reduces when updating conforms to Bayes' Rule.

In common with [5], the present paper adapts the Gul and Pesendorfer [9, 10] model of temptation and self-control. 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 $(s_{t+2}, s_{t+3}, \dots)$. 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, and thus retroactively change her prior. Then she applies Bayes' Rule to the new prior. The resulting posterior belief 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. The implication for behavior is the urge to choose current consumption so as to maximize expected utility conditioning on the new prior as opposed to the initial prior. Thus temptation refers to giving in to one’s urges, which here stem from a change in beliefs. Temptation might be resisted but at a cost.

GP show that temptation and self-control are revealed through preference over menus. Menus play a central role in [5] and in this model as well. We model preferences over *contingent menus* and show that these preferences reveal behavior that is consistent with non-Bayesian updating rules.

The paper proceeds as follows: Section 2 defines the formal domain of choice, the space of contingent menus, and then the functional form for utility. Section 3 provides axiomatic foundations. Section 4 illustrates the nature and scope of the model by describing axiomatic specializations that capture specific updating biases.³ Section 5 specializes further to capture an agent who is trying to learn about parameters as in the Bayesian model with an exchangeable prior. Some results are provided concerning what is learned in the long run. Section 6 concludes. Proofs are collected in appendices.

2. UTILITY

2.1. Primitives

The model’s primitives include:

- time $t = 0, 1, 2, \dots$
- (finite) period state space S
full state space is $\Pi_{t=1}^{\infty} S_t, S_t = S$ for all t
- period consumption space $C_t = C$
compact metric and a mixture space

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.

³Readers who are more interested in the functional forms implied by our model than in their axiomatic foundations may wish to skip Sections 3 and 4 and proceed directly to Section 5; the latter is in large part self-contained.

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 [1, 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

$$F : S \longrightarrow \mathcal{M}(C \times C), \quad (2.1)$$

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

$$\mathcal{C} \underset{\text{homeo}}{\approx} (\mathcal{M}(C \times C))^S. \quad (2.2)$$

Hence, we identify any element of \mathcal{C} with a mapping F as in (2.1).


We study preferences on \mathcal{C} . In fact, since the model is dynamic and choices are made at each history $s_1^t = (s_1, \dots, s_t)$, we take as primitive a process of preference relations $(\succeq_t)_{t=0}^\infty$, where \succeq_t is the order on \mathcal{C} prevailing at t ; \succeq_t depends on the history s_1^t , but this dependence is suppressed in the notation.

Though the domain \mathcal{C} is time stationary, that is, the objects of choice at each t are elements of the same set \mathcal{C} , when we wish to emphasize that a particular choice is made at t , we write that the agent chooses contingent menu $F_t \in \mathcal{C}_t$,

$$F_t : S_{t+1} \longrightarrow \mathcal{M}_{t+1} \equiv \mathcal{M}(C_{t+1} \times C_{t+1}), \quad (2.3)$$

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 time line is as follows:



At t^+ , the agent chooses a contingent menu F_t in \mathcal{C}_t . She does this as though anticipating the following: at $(t+1)^-$ a signal s_{t+1} is realized; and at $t+1$, she updates and chooses some (c_{t+1}, F_{t+1}) from the menu $F_t(s_{t+1})$. This leaves her with F_{t+1} at $(t+1)^+$, which she may want to replace if possible. If she does not, and she will not in our model, the single utility function at time 0^+ determines F_0 and subsequently the process of consumption-contingent-menu pairs over the entire horizon, in parallel with the standard Bayesian model.

The preceding must be qualified, however. Our formal model addresses exclusively the *choice of contingent menus* at each t . The needed model of *choice out of menus* after realization of each period's signal is 'suggested' but is not part of the formal setup. Similar gaps in foundations arise in both GP's papers and in Epstein's three-period framework. In [9], GP show how this gap may be filled by the study of suitably extended preferences and a similar solution could be provided here. Note that foundations provided in this way are subject to the difficulty pointed out in [9, p. 1415], namely the lack of a revealed preference basis for extended preferences.

Subclasses of \mathcal{C} that provide complete or partial commitment are of special interest. The contingent menu F provides commitment for the next period if $F(s)$ is a singleton for each s . The set $\mathcal{C}^0 \subset \mathcal{C}$ of perfect (for all future periods) commitment prospects, is defined by:

$$F \in \mathcal{C}^0 \iff \forall s \in S, \exists (c', F') \in C \times \mathcal{C}^0 \text{ s.t. } F(s) = \{(c', F')\}. \quad (2.4)$$

Repeated application of (2.4) implies that each F in \mathcal{C}^0 determines a unique (random variable) consumption process $c^F = (c_t^F)$.

Let \mathcal{C}_{+1}^0 consist of the subset of \mathcal{C}^0 for which there is not only commitment, but where also all relevant uncertainty is resolved in the next period. An example

is a (one-step-ahead) *bet on the event* $G \subset S$, which pays off with a good deterministic consumption stream if the state next period lies in G and with a poor one otherwise.

Any consumption stream can be viewed as a special contingent menu. We refer to such streams as risky, because each c_t is a lottery over period t consumption. We identify each risky stream, denoted \vec{c} , with an element of \mathcal{C}^r ($\mathcal{C}^r \subset \mathcal{C}_{+1}^0 \subset \mathcal{C}^0 \subset \mathcal{C}$); (c, \vec{c}) denotes the obvious consumption stream.

2.2. Functional Form

We describe the functional form for $\mathcal{U}_t : \mathcal{C}_t \rightarrow \mathbb{R}^1$ representing \succeq_t .⁴ Though it is likely unfamiliar and may seem complicated, the functional form is *not* ad hoc. Section 3 provides the axiomatic underpinnings.

Components of the functional form include: a discount factor $0 < \delta < 1$, $u : C \rightarrow \mathbb{R}^1$ linear, continuous and nonconstant, $p_0 \in \Delta(S_1)$ and the adapted process $(p_t, q_t, \alpha_t)_{t \geq 1}$, where⁵

$$\alpha_t \in \mathbb{R}_+^1, \quad p_t, q_t \in \Delta(S_{t+1}), \quad q_t \ll p_t, \quad \text{for } t \geq 1.$$

It is convenient to define the measure m_t on S_{t+1} by

$$m_0 = p_0 \quad \text{and} \quad m_t = \frac{p_t + \alpha_t q_t}{1 + \alpha_t}, \quad t \geq 1. \quad (2.5)$$

Utilities are given (for $t \geq 0$) by

$$\mathcal{U}_t(F_t) = \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dm_t(s_{t+1}), \quad F_t \in \mathcal{C}_t, \quad (2.6)$$

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

$$U_{t+1}(M_{t+1}, s_{t+1}) = \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} (1 + \alpha_{t+1}) \left\{ u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) d \left(\frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}} \right) \right\} \quad (2.7)$$

⁴That is, $\mathcal{U}_t(\cdot; s_1^t)$ is the utility function. Similarly, below we often write $p_t(\cdot)$ rather than $p_t(\cdot | s_1^t)$. When we want to emphasize dependence on the last observation s_t , we write $p_t(\cdot | s_t)$.

⁵ $\Delta(S)$ is the set of probability measures on the finite set S . For any two measures p and q in $\Delta(S)$, write $q \ll p$ if q is absolutely continuous with respect to p . A stochastic process (X_t) on $\Pi_1^\infty S_\tau$ 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 $\{s_1\} \times \dots \times \{s_t\} \times \Pi_{t+1}^\infty S_\tau$.

$$- \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \alpha_{t+1} \left\{ u(c'_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F'_{t+1}(s_{t+2}), s_{t+2}) dq_{t+1}(s_{t+2}) \right\}.$$

The Bayesian intertemporal utility model (1.1) is specified by u , δ and the process of one-step-ahead conditionals, which determines a unique prior on the full state space $\Pi_{t=1}^{\infty} S_t$. It is obtained as the special case where $\alpha_t(q_t - p_t) \equiv 0$ for all t . Then, (2.7) reduces to

$$U_{t+1}(M_{t+1}, s_{t+1}) = \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dp_{t+1} \right\}.$$

This is the standard model in the sense that it extends the model of utility over consumption processes given by (1.1) to contingent menus by assuming that menus are valued according to the best alternative they contain (a property termed *strategic rationality* by Kreps [12]). In particular, time t conditional beliefs about the future are obtained by applying Bayes' Rule to the measure on $\Pi_1^{\infty} S_{\tau}$ that is induced by the one-step-ahead conditionals $(p_{\tau})_1^{\infty}$.

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 4 and 5 provide several examples. See also [5] for discussion in the context of a three-period model.

In terms of the time line described above, at t^+ the agent chooses a contingent menu F_t in \mathcal{C} according to the utility function \mathcal{U}_t . At $(t+1)^-$ a signal s_{t+1} is realized. Then she updates and chooses some (c_{t+1}, F_{t+1}) from the menu $F_t(s_{t+1})$. The functional form suggests that choice from the menu is made by solving⁶

$$\max_{(c_{t+1}, F_{t+1}) \in F_t(s_{t+1})} \left\{ u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dm_{t+1}(s_{t+2}) \right\}. \quad (2.8)$$

This leaves her with $F_{t+1}(s_{t+1})$ at $(t+1)^+$ and so on. Note that dynamic inconsistency does not arise. For example, consider the time t^+ choice F_t and whether replanning at $(t+1)^-$ will lead to deviations. However, at that point, after s_{t+1} has been realized and before updating, $U_{t+1}(\cdot, s_{t+1})$ represents preference over menus. Dynamic consistency between \mathcal{U}_t and the collection of utilities $\{U_{t+1}(\cdot, s_{t+1}) : s_{t+1} \in S_{t+1}\}$ is evident from (2.6). Thus at $t+1$, when she updates,

⁶We adopt the suggested model of choice out of menus for purposes of interpretation and intuition. As noted above, formal justification can be provided.

the agent is left to choose *out of the menu* $F_t(s_{t+1})$. This choice cannot conflict with prior choices even in principle, because choice *out of menus* is of a different nature than choice *of menus* - this is the insight of GP. For the same reason, the resulting choice F_{t+1} cannot be overturned by \mathcal{U}_{t+1} . The bottom line is that dynamic choice behavior can be viewed as resulting from optimizing the single utility function \mathcal{U}_0 , and also admits a recursive characterization corresponding to the repeated solution of (2.8).

2.3. Interpretation

Turn to interpretation of the model (2.5)-(2.7). First, for a contingent menu F_t that provides commitment ($F_t \in \mathcal{C}^0$), compute that

$$\begin{aligned} \mathcal{U}_t(F_t) &= \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) \, dm_t(s_{t+1}) \\ &= \int_{S_{t+1} \times \dots} \sum_{\tau=t+1}^{\infty} \delta^{\tau-t-1} u(c_{\tau}^{F_t}) \, dP_t(\cdot | s_1^t), \end{aligned} \quad (2.9)$$

where c^{F_t} is the consumption process induced by F_t as explained following (2.4), and where $P_t(\cdot | s_1^t)$ is the unique measure on $\Pi_{t+1}^{\infty} S_{\tau}$ satisfying (for each T)

$$P_t(s_{t+1}, s_{t+2}, \dots, s_T | s_1^t) = m_t(s_{t+1}) p_{t+1}(s_{t+2} | s_1^{t+1}) \times \dots \times p_{T-1}(s_T | s_1^{T-1}). \quad (2.10)$$

Thus \succeq_t restricted to \mathcal{C}^0 conforms to subjective expected (intertemporally additive) utility with prior P_t , which we therefore refer to as the *commitment prior*. When \succeq_t is further restricted to \mathcal{C}_{+1}^0 , (all uncertainty resolves at $t+1$), it has prior m_t ; for example, m_t represents the time t 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 preceding provides behavioral meaning for the measures m_t and p_t . It also clarifies one sense in which updating is non-Bayesian: P_{t+1} is not the Bayesian update of P_t . In other words, commitment beliefs at $t+1$ are *not* derived from Bayesian updating of the commitment prior at t .

For further interpretation, note that (2.7) can be written in the form:

$$U_{t+1}(M_{t+1}, s_{t+1}) = \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \{ U_{t+1}^{comm}(c_{t+1}, F_{t+1}) + \quad (2.11)$$

$$\alpha_{t+1} \left[V_{t+1}(c_{t+1}, F_{t+1}) - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} V_{t+1}(c'_{t+1}, F'_{t+1}) \right],$$

where U_{t+1}^{comm} gives utility using p_{t+1} , the Bayesian update of the time t commitment prior P_t , that is,

$$U_{t+1}^{comm}(c_{t+1}, F_{t+1}) = u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dp_{t+1}(s_{t+2}), \quad (2.12)$$

and where V_{t+1} is an expected utility function using the measure q_{t+1} , that is,

$$V_{t+1}(c_{t+1}, F_{t+1}) = u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) dq_{t+1}(s_{t+2}). \quad (2.13)$$

The time t perspective calls for choosing *out of the menu* M_{t+1} by maximizing U_{t+1}^{comm} . However, after seeing the realization s_{t+1} , the agent revises her view of what it implies about the future and adopts the conditional one-step-ahead belief given by q_{t+1} . Thus she is *tempted* to maximize V_{t+1} . To the extent that she resists this temptation and chooses (c_{t+1}, F_{t+1}) , she incurs the (utility) self-control cost

$$\alpha_{t+1} \left[V_{t+1}(c_{t+1}, F_{t+1}) - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} V_{t+1}(c'_{t+1}, F'_{t+1}) \right];$$

thus α_{t+1} parametrizes the cost of self-control. Finally, (2.11) suggests that choice out of the menu M_{t+1} is made *as though* she adopts the *compromise one-step-ahead belief* $m_{t+1} = \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}}$. The use of m_{t+1} to guide choice out of the menu, rather than p_{t+1} , the one-step-ahead Bayesian conditional of P_t , is another behavioral expression of non-Bayesian updating.⁷

3. AXIOMATIC FOUNDATIONS

Consider axioms for the preference process (\succeq_t) , where each \succeq_t is defined on \mathcal{C} . We adapt the axioms used in the three-period setting of [5], and then add to them assumptions that are specific to the infinite horizon setting. The former axioms are ‘static’ in that they deal with preferences at each given time, while axioms in the second group are ‘dynamic’ in that they relate preferences at different times.

⁷We remind the reader of the option of skipping directly to Section 5.

3.1. A First Group of Axioms

For most of these axioms, intuition is similar to that provided in [5] and thus the discussion is brief.

The first three axioms are standard.⁸

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

Axiom 2 (Continuity). Both $\{F \in \mathcal{C} : F \succeq_t G\}$ and $\{F \in \mathcal{C} : G \succeq_t F\}$ are closed.

Axiom 3 (Nondegeneracy). $F \succ_t G$ for some F and G in \mathcal{C} .

In Section 2.1, we described a way to mix any two elements in \mathcal{C} . Thus we can state the Independence axiom appropriate for our setting.

Axiom 4 (Independence). For every $0 < \lambda \leq 1$, $F \succeq_t G$ if and only if $\lambda F + (1 - \lambda) F' \succeq_t \lambda G + (1 - \lambda) F'$.

Define the union of contingent menus statewise, that is,

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

The counterpart of GP's central axiom is:

Axiom 5 (Set-Betweenness). For all states s and all menus F and G in \mathcal{C} such that $G(s') = F(s')$ for all $s' \neq s$,

$$F \succeq_t G \implies F \succeq_t F \cup G \succeq_t G. \quad (3.1)$$

The axioms Order, Continuity and Independence imply weak separability across states, so that the rankings appearing in (3.1) are independent of the common outcomes in states $s' \neq s$. Thus one can interpret them as being rankings that are conditional on the realization of s but before updating has been done; this is the perspective of $(t + 1)^-$ in the time line drawn earlier. For example, the hypothesis $F \succeq_t G$ means that after seeing s , the agent strictly prefers to have the menu $F(s)$ rather than $G(s)$ from which to choose at $(t + 1)$ after updating. Similar interpretations in terms of conditional preference could be adopted in many instances below.

⁸States s vary over S . Unless otherwise specified, time t varies over $1, 2, \dots$

Set-Betweenness allows for both temptation and self-control - GP show this in their setting and [5] adapts their interpretation to the domain of (three-period) contingent menus. Interpret $F \succ_t F \cup G$ as a situation where (conditional on s) the decision-maker is tempted by G . She succumbs to this temptation if also $F \cup G \sim_t G$, while she is able to exert self-control and resist if $F \cup G \succ_t G$. For perspective, note that temptations do not exist for a standard decision-maker who evaluates a menu by its best element. She 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 [12, Ch. 13], we call this axiom *strategic rationality*.

We model an agent who faces temptations because of “changing beliefs” at the updating stage. This connection between temptation and updating is imposed largely through the next two axioms that weaken strategic rationality. To express them, we need some additional notation. Call the contingent menu ℓ constant if $\ell(s') = \ell(s)$ for all s' and s ; each such ℓ can be identified with its range, an element of $\mathcal{M}(C \times C)$. For any state s , $F \in \mathcal{C}$ and $\ell \in \mathcal{M}(C \times C)$, $\ell s F$ denotes the contingent menu given by

$$(\ell s F)(s') = \begin{cases} F(s') & \text{if } s' \neq s \\ \ell & \text{otherwise.} \end{cases}$$

Similarly, for any $M \subset \mathcal{C}$ and $L \subset \mathcal{M}(C \times C)$,

$$LsM = \{\ell s F : \ell \in L, F \in M\} \subset \mathcal{C}. \quad (3.2)$$

Sets of contingent menus having this form have an important property: imagine having to make a choice out of LsM . Typically, the time at which this choice must be made would be important - one would be better off if it could be made after knowing whether s is true. However, the menu LsM is sufficiently rich so that any ex post choice can be replicated even if the choice must be made before knowing if s is true: for example, if $\ell^* \in L$ would be chosen ex post if s is realized and if $F^* \in M$ would be chosen otherwise, then the ex ante choice of $\ell^* s F^* \in LsM$ would produce the identical contingent plan. As a result, we are free to think of choice out of LsM as occurring after learning if s is realized.

For any $M \subset \mathcal{C}$ and $c \in C$, denote by (c, M) the set

$$(c, M) \equiv \{c\} \times M \subset C \times \mathcal{C};$$

similarly for the meaning of (K, M) for any $K \subset C$. When M is a singleton $\{F'\}$, we write simply (K, F') rather than $(K, \{F'\})$. When $L \subset \mathcal{M}(C \times C)$, interpret (K, L) as above except that L is identified with a set of constant contingent menus. Finally, for any F in \mathcal{C} , $(F_{-s_{t+1}}, (c, Ls_{t+2}M))$ denotes the contingent menu that delivers $F(s'_{t+1})$ if $s'_{t+1} \neq s_{t+1}$ and $\{c\} \times Ls_{t+2}M$ otherwise.

Now we can state:

Axiom 6 (Restricted Strategic Rationality (RSR)). *For all (s_{t+1}, s_{t+2}) , c , contingent menus F , menus $M \subset \mathcal{C}$ and $L', L \subset \mathcal{M}(C \times C)$,*

$$\begin{aligned} (F_{-s_{t+1}}, (c, L's_{t+2}M)) \succ_t (F_{-s_{t+1}}, (c, Ls_{t+2}M)) &\implies \\ (F_{-s_{t+1}}, (c, L's_{t+2}M)) \sim_t (F_{-s_{t+1}}, (c, (L' \cup L)s_{t+2}M)). \end{aligned}$$

Interpret the indicated rankings as conditional on having observed s_{t+1} but before updating. Then, as above, the hypothesized ranking indicates the (conditional) preference to receive L' rather than L in the state s_{t+2} . The point is that updating is not relevant to this ranking, because the comparison is between two menus that differ only in the single state s_{t+2} . Such a comparison does not involve trade-offs across states and hence does not depend on beliefs about S_{t+2} or on updating in response to s_{t+1} . Because temptations arise only with the change in beliefs that occurs when updating, a form of strategic rationality obtains for such comparisons. This explains the implied indifference (conditional on s_{t+1}) between receiving L' or $L' \cup L$ in state s_{t+2} .

The next axiom also reflects the connection between temptation and updating.

Axiom 7 (Consumption Strategic Rationality (CSR)). *For all s_{t+1} , contingent menus F and F' , and all consumption menus $K', K \subset C$,*

$$\begin{aligned} (F_{-s_{t+1}}, (K', F')) \succeq_t (F_{-s_{t+1}}, (K, F')) &\implies \\ (F_{-s_{t+1}}, (K', F')) \succeq_t (F_{-s_{t+1}}, ((K' \cup K), F')). \end{aligned}$$

Conditional on s_{t+1} but before updating, the agent prefers (K', F') to (K, F') , which differ only in the menus (K' versus K) for consumption at $t + 1$. After choosing one of these and then updating in response to s_{t+1} , the agent selects from the chosen menu. However, beliefs about S_{t+2} are not directly relevant to the choice of consumption at $t + 1$. Beliefs could be indirectly relevant because they influence the evaluation of the options F' for the future, but this connection

is mute if the ranking of current consumption menus is weakly separable from the future. In that case, updating is not relevant to the ranking hypothesized in the axiom and the indicated form of strategic rationality applies.

Given any contingent menu F in \mathcal{C} and event $A \subset S$, denote by $proj_A F$ the restriction of F to A . Similarly, if M is a subset of \mathcal{C} , denote the set of restrictions (or projections) by

$$proj_A M = \{proj_A F : F \in M\}.$$

When $A = \{s\}$, write simply $proj_s F$ or $proj_s M$.

Some of the following axioms make use of a notion of nullity that we now define. For any $E \subset S_{t+2}$, say that (s_{t+1}, E) is \succeq_t -null if $G \sim_t F$ for all G and F satisfying (for some c)

$$G(s'_{t+1}) = F(s'_{t+1}) \quad \text{for all } s'_{t+1} \neq s_{t+1}, \quad (3.3)$$

$$\begin{aligned} G(s_{t+1}) &= (c, M^G), \quad F(s_{t+1}) = (c, M^F), \\ proj_{S_{t+2} \setminus E} M^G &= proj_{S_{t+2} \setminus E} M^F. \end{aligned} \quad (3.4)$$

Because G and F agree for $s'_{t+1} \neq s_{t+1}$ and induce the same set of restrictions on $S_{t+2} \setminus E$, then G and F “differ only on $\{s_{t+1}\} \times E$.” The latter being null presumably means that any such G and F are indifferent. When $E = \{s_{t+2}\}$ is a singleton, refer to nullity of (s_{t+1}, s_{t+2}) rather than of $(s_{t+1}, \{s_{t+2}\})$. Say that s_{t+1} is \succeq_t -null if (s_{t+1}, S_{t+2}) is \succeq_t -null as just defined, that is, if $G \sim_t F$ whenever (3.3) is satisfied.

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

Axiom 8 (State Independence). For all (s_{t+1}, s_{t+2}) and \succeq_t -non-null (s'_{t+1}, s'_{t+2}) , and for all c , contingent menus F , menus $M \subset \mathcal{C}$ and $L', L \subset \mathcal{M}(C \times C^r)$,

$$\begin{aligned} \left(F_{-s'_{t+1}}, (c, L' s'_{t+2} M) \right) &\succeq_t \left(F_{-s'_{t+1}}, (c, L s'_{t+2} M) \right) \implies \\ \left(F_{-s_{t+1}}, (c, L' s_{t+2} M) \right) &\succeq_t \left(F_{-s_{t+1}}, (c, L s_{t+2} M) \right). \end{aligned}$$

The two contingent menus in the hypothesized ranking differ only through the difference between L' and L . Thus the indicated ranking states roughly that the

agent prefers to receive L' rather than L in the state (s'_{t+1}, s'_{t+2}) .⁹ Suppose that risk attitudes are not state-dependent. The difference between (s'_{t+1}, s'_{t+2}) and (s_{t+1}, s_{t+2}) may still matter because the two states may imply different beliefs about the future and hence for the payoffs from choosing out of L' rather than out of L . However, this is not the case when L' and L are menus of risky prospects, that is, $L', L \subset \mathcal{M}(C \times \mathcal{C}^r)$. Then the preference for L' should prevail also at (s_{t+1}, s_{t+2}) .

Axiom 9 (Absolute Continuity). *For all (s_{t+1}, s_{t+2}) , c in C , and for all contingent menus F satisfying $F(s_{t+1}) = (c, M)$, if*

$$(F_{-s_{t+1}}, (c, L's_{t+2}M)) \sim_t (F_{-s_{t+1}}, (c, Ls_{t+2}M)) \quad (3.5)$$

for all $L', L \subset \mathcal{M}(C \times \mathcal{C})$, then

$$F' = (F_{-s_{t+1}}, (c, Ls_{t+2}M)) \sim_t F, \quad \text{where } L = \text{proj}_{s_{t+1}}M. \quad (3.6)$$

Note that, given the hypothesis, the conclusion (3.6) is equivalent to asserting that (s_{t+1}, s_{t+2}) is \succeq_t -null.

The two contingent menus appearing in (3.5) differ only in state (s_{t+1}, s_{t+2}) , where they offer either L' or L . This difference is a matter of indifference, regardless of the nature of L' and L , which suggests that from the perspective of t , (s_{t+1}, s_{t+2}) is viewed as impossible. Suppose that the agent anticipates that she will continue to view s_{t+2} as impossible also if s_{t+1} is realized at $t + 1$ (roughly, that conditional beliefs are absolutely continuous with respect to ex ante beliefs). Then she anticipates that given s_{t+1} and facing M , she will believe that s_{t+2} is impossible and thus will find herself ‘as if’ she can choose from M after learning if s_{t+2} is true. But, as explained in the discussion surrounding (3.2), this is precisely the situation she would anticipate if she expected that state s_{t+1} would lead to the menu $Ls_{t+2}M$, where $L = \text{proj}_{s_{t+1}}M$. Conclude that at t she will be indifferent between F and $(F_{-s_{t+1}}, (c, Ls_{t+2}M))$.

The last two axioms, like Consumption Strategic Rationality above, do not have counterparts in [5] because consumption was assumed there to be limited to the terminal time. These axioms are needed, however, in any multi-period setup once intermediate consumption is permitted as it is here.

⁹The earlier point about equivalence between ex ante and ex post choice out of $Ls'_{t+2}M$ is important here.

Axiom 10 (Risk Preference). *There exist $0 < \delta < 1$ and $u : C \rightarrow \mathbb{R}^1$ non-constant, linear and continuous, such that, for each F in \mathcal{C} , \vec{d} and \vec{e} in \mathcal{C}^r , and for every t and s_{t+1} such that s_{t+1} is \succeq_t -non-null,*

$$\begin{aligned} (F_{-s_{t+1}}, \vec{d}) \succeq_t (F_{-s_{t+1}}, \vec{e}) &\iff \\ \Sigma_{t+1} \delta^{\tau-(t+1)} u(\vec{d}_\tau) &\geq \Sigma_{t+1} \delta^{\tau-(t+1)} u(\vec{e}_\tau). \end{aligned} \quad (3.7)$$

The axiomatic characterization of the utility function over streams of lotteries appearing in (3.7) is well known (see [4], for example). Because risk preferences are not our primary focus, we content ourselves with the statement of the above unorthodox ‘axiom.’

Axiom 11 (Invariant Discounting). *For all s_{t+1} , all c, d and e in C , \vec{d} and \vec{e} in \mathcal{C}^r , and F in \mathcal{C} , if:*

$$(F_{-s_{t+1}}, (c, F')) \succ_t (F_{-s_{t+1}}, \{(c, F'), (d, \vec{d})\}) \succ_t (F_{-s_{t+1}}, (d, \vec{d})) \quad (3.8)$$

and similarly with (e, \vec{e}) in place of (d, \vec{d}) , then

$$\begin{aligned} (F_{-s_{t+1}}, \{(c, F'), (e, \vec{e})\}) \succeq_t (F_{-s_{t+1}}, \{(c, F'), (d, \vec{d})\}) &\quad (3.9) \\ \iff (F_{-s_{t+1}}, (d, \vec{d})) \succeq_t (F_{-s_{t+1}}, (e, \vec{e})). \end{aligned}$$

Conditional on s_{t+1} , each of (d, \vec{d}) and (e, \vec{e}) tempts (c, F') and in each case the temptation is resisted. The question is which is more tempting. Because (d, \vec{d}) and (e, \vec{e}) correspond to constant contingent menus, they are purely risky. Thus their comparison reflects the different trade-offs that they offer between times $t + 1$ and the future, that is, it reflects discounting. This is so both in evaluating which is more tempting and also when comparing them as commitment prospects. Accordingly, (3.9) says that the agent discounts in the same way when deciding how tempting are (d, \vec{d}) and (e, \vec{e}) as she does when ranking them under commitment.

3.2. ‘Dynamic’ Axioms

All of the above axioms have counterparts in a three-period setting (with intermediate consumption) where updating is done only once. The remaining axioms impose structure that is specific to a multi-period setting. Specifically, they relate preferences at different times.¹⁰

Axiom 12 (Restricted Recursivity (RRC)). For all \succeq_t -non-null (s_{t+1}, s_{t+2}) , all c , contingent menus $F, M \subset \mathcal{C}$ and $\ell', \ell \subset \mathcal{C} \times \mathcal{C}$,

$$\begin{aligned} (F_{-s_{t+1}}, (c, \{\ell'\}_{s_{t+2}}M)) &\succeq_t (F_{-s_{t+1}}, (c, \{\ell'\}_{s_{t+2}}M)) \\ \iff (F_{-s_{t+2}}, (c, \ell')) &\succeq_{t+1} (F_{-s_{t+2}}, (c, \ell)). \end{aligned}$$

Suppose that at t the agent, looking forward two periods, prefers to receive ℓ' in (s_{t+1}, s_{t+2}) to receiving ℓ . Then also at $t+1$, after s_{t+1} is realized, if she looks forward one period, she prefers to receive ℓ' in s_{t+2} to receiving ℓ . Similarly for the converse.

Axiom 13 (Bias Persistence). For all s_{t+1} and c , and for all F, F', G and H in \mathcal{C} : if **either**

$$(F_{-s_{t+1}}, (c, F')) \sim_t (F_{-s_{t+1}}, \{(c, F'), (c, G)\}) \succ_t (F_{-s_{t+1}}, (c, G)), \quad (3.10)$$

or

$$\begin{aligned} (F_{-s_{t+1}}, (c, F')) &\succ_t (F_{-s_{t+1}}, \{(c, F'), (c, H)\}) \succ_t (F_{-s_{t+1}}, (c, H)), \\ (F_{-s_{t+1}}, (c, G)) &\succ_t (F_{-s_{t+1}}, \{(c, G), (c, H)\}) \succ_t (F_{-s_{t+1}}, (c, H)), \\ \text{and } (F_{-s_{t+1}}, \{(c, F'), (c, H)\}) &\succeq_t (F_{-s_{t+1}}, \{(c, G), (c, H)\}), \end{aligned} \quad (3.11)$$

then

$$F' \succeq_{t+1} G. \quad (3.12)$$

¹⁰In the next axiom, \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. In addition, it may be clarifying to include the time subscripts and write $\ell', \ell \subset \mathcal{C}_{t+2} \times \mathcal{C}_{t+2}$.

Consider the rankings hypothesized in (3.11). They indicate that ex ante at t , prior to the realization of s_{t+1} and in anticipation thereof, the agent anticipates that H will be tempting given either G or F' . (Because c is fixed throughout, it is suppressed in this explication). Further, she anticipates having self-control at both $\{G, H\}$ and at $\{F', H\}$, and choosing G from $\{G, H\}$ and F' from $\{F', H\}$. The final ranking in (3.11) indicates that she anticipates H being less tempting to F' than to G , and consequently that she will prefer to choose F' from $\{F', H\}$ rather than G from $\{G, H\}$. The axiom then requires that at $t + 1$, *after* s_{t+1} is in fact realized, she should also prefer F' to G .

The connection to persistence in updating is implicit. Updating, whether anticipated or actual, underlies both the rankings in (3.11) and that in (3.12). Thus the axiom can be interpreted as requiring that the updating anticipated to underlie the choice out of menus, say of F' from $\{F', H\}$ and G from $\{G, H\}$, persists and is used at $t + 1$ when ranking contingent menus. Though this requirement may at first glance seem necessary for a dynamically coherent model, this is not the case because the two choice problems indicated are different: the choice at $t + 1$ is between the individual contingent menus F' and G , while the problem anticipated at t involves choice from specific menus containing the tempting alternative H .

The interpretation differs slightly when considering the alternative hypothesis (3.10). The latter indicates anticipation that G will not tempt F' on realization of s_{t+1} . Thus she anticipates that if s_{t+1} is realized, then she will choose F' from $\{F', G\}$. According to (3.12), she should then prefer F' to G at $t + 1$ *after* s_{t+1} is realized.

The role of the axiom is to deliver the relation (2.5) between the various measures m_t , p_t and q_t .

3.3. Representation Result

We are finally in a position to state our first result.

Theorem 3.1. *If the process (\succeq_t) of preferences satisfies axioms (1)-(13), then it admits representation of the form (2.5)-(2.7), where u , δ , p_0 and $(\alpha_t, p_t, q_t)_{t \geq 1}$ satisfy the properties stated there.*

Conversely, suppose that equations (2.5)-(2.7) admit a unique solution (U_t) , where $U_t(\cdot, s_1^t) : \mathcal{M}(C \times C) \rightarrow \mathbb{R}^1$ is continuous and uniformly bounded in the sense that

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

Let \succeq_t be represented by $\mathcal{U}_t(\cdot, s_1^t) : \mathcal{C} \longrightarrow \mathbb{R}^1$,

$$\mathcal{U}_t(F; s_1^t) = \int_{S_{t+1}} U_{t+1}(F(s_{t+1}), s_1^{t+1}) dm_t(s_{t+1}), \quad F \in \mathcal{C}.$$

Then (\succeq_t) satisfies axioms (1)-(13). Finally, a sufficient condition for the existence of a unique solution (U_t) as above is that

$$\left(1 + 2 \sup_{t, s_1^t} |\alpha_t(s_1^t)|\right) \delta < 1. \quad (3.13)$$

Remark 1. The proof reveals that if Bias Persistence is dropped, then the remaining axioms imply the representation (2.6)-(2.7) where the relation (2.5) between m_t and the other measures is weakened to the absolute continuity requirement $p_t \ll m_t$.

We conclude with an examination of the uniqueness properties of the representation.

Corollary 3.2. Let (\succeq_t) satisfying the axioms admit representations by both $(\delta, u, p_0, (\alpha_t, p_t, q_t)_{t \geq 1})$ and $(\delta', u', p'_0, (\alpha'_t, p'_t, q'_t)_{t \geq 1})$ as assured by the theorem. Then $\delta' = \delta$, $u' = au + b$ for some $a > 0$, and

$$p'_t = p_t, \quad m'_{t+1} = \frac{p'_{t+1} + \alpha'_{t+1}q'_{t+1}}{1 + \alpha'_{t+1}} = \frac{p_{t+1} + \alpha_{t+1}q_{t+1}}{1 + \alpha_{t+1}} = m_{t+1} \quad \text{for } t \geq 0. \quad (3.14)$$

Moreover, if $t \geq 0$ and s_{t+1} are such that

$$\begin{aligned} (F_{-s_{t+1}}, (c, M')) &\succeq_t (F_{-s_{t+1}}, (c, M)) \quad \text{and} \\ (F_{-s_{t+1}}, (c, M')) &\approx_t (F_{-s_{t+1}}, (c, M' \cup M)) \end{aligned} \quad (3.15)$$

for some $c \in \mathcal{C}$ and $M', M \subset \mathcal{C}$, then

$$(\alpha'_{t+1}(s_{t+1}), q'_{t+1}(\cdot | s_{t+1})) = (\alpha_{t+1}(s_{t+1}), q_{t+1}(\cdot | s_{t+1})). \quad (3.16)$$

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 (3.15). Once strategic rationality is excluded, the strong uniqueness property in (3.16) is valid.¹¹

4. SOME SPECIFIC UPDATING BIASES

The framework described in Theorem 3.1 is rich. One way to see this is to focus on one-step-ahead beliefs at any time $t + 1$. As pointed out in Section 2.3, these are represented by $m_{t+1} = \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}}$, while Bayesian updating of the time t commitment prior 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.

In this section, we go further and describe axiomatic 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. The axioms imply restrictions 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.

4.1. Prior-Bias

The agent at $t+1$, when updating in response to seeing s_{t+1} , may attach inordinate weight to his prior view of S_{t+2} , that is, to his view at t . To express this, let $(c_{t+1}, G) \in C \times \mathcal{C}$, where c_{t+1} is consumption for $t + 1$ and where G represents the random future beyond $t + 1$; formally $G : S_{t+2} \rightarrow \mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$. Denote by $\overline{(c_{t+1}, G)}$ the contingent menu in \mathcal{C} that assigns the singleton $\{(c_{t+1}, G)\}$ to every s_{t+1} . The ranking of such contingent menus induced by \succeq_t reflects the noted prior view at t .

¹¹The proof is analogous to the proof of the uniqueness properties in [5]. For example, (3.14) follows from the SEU representation of preference restricted to commitment prospects \mathcal{C}^0 (recall (2.9)-(2.10)) and uniqueness of the representing prior in an Anscombe-Aumann setting.

Axiom 14 (Prior-Bias). For all s_{t+1} and c_{t+1} , and all F' , F and G in \mathcal{C} : if

$$(F'_{-s_{t+1}}, (c_{t+1}, F)) \succ_t (F'_{-s_{t+1}}, (c_{t+1}, G)) \text{ and} \quad (4.1)$$

$$\overline{(c_{t+1}, F)} \sim_t \overline{(c_{t+1}, G)}, \quad (4.2)$$

then

$$(F'_{-s_{t+1}}, (c_{t+1}, F)) \sim_t (F'_{-s_{t+1}}, \{(c_{t+1}, F), (c_{t+1}, G)\}). \quad (4.3)$$

To interpret the axiom, we suppress the fixed consumption c_{t+1} (and do the same for interpretations in the sequel). Condition (4.1) states that, conditionally on s_{t+1} , the agent strictly prefers to commit to F rather than to G . According to (4.2), she is indifferent between them that ex ante at t . Under these circumstances, she is not tempted by G conditionally on s_{t+1} . Thus the absence 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 realization of s_{t+1} . This indicates excessive influence of prior beliefs at the updating stage.

Prior-Bias begs the question what happens to temptation if the indifference in (4.2) is not satisfied? We consider two alternative strengthenings of the axiom that provide different answers.

Label by **Positive Prior-Bias** the axiom obtained when (4.2) is replaced by

$$\overline{(c_{t+1}, F)} \succeq_t \overline{(c_{t+1}, G)}. \quad (4.4)$$

Then 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 (4.2) is replaced by

$$\overline{(c_{t+1}, F)} \preceq_t \overline{(c_{t+1}, G)}. \quad (4.5)$$

In this case, G is preferred ex ante 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 4.1. Let (\succeq_t) satisfy the axioms in Theorem 3.1. Then it satisfies Prior-Bias if and only if it admits representation as in the Theorem where in addition

$$q_{t+1}(\cdot | s_{t+1}) = (1 - \lambda_{t+1})p_{t+1}(\cdot | s_{t+1}) + \lambda_{t+1} \left[\sum_{s'_{t+1}} m_t(s'_{t+1}) p_{t+1}(\cdot | s'_{t+1}) \right], \quad (4.6)$$

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 (4.6) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$.

One can see from (2.10) that the measure $\sum_{s'_{t+1}} m_t(s'_{t+1}) p_{t+1}(\cdot | s'_{t+1})$ represents commitment beliefs about S_{t+2} held at t ; we refer to it as the (ex ante) view of S_{t+2} held at t . Thus Prior-Bias is characterized by q_{t+1} being expressible as a linear combination of p_{t+1} (the Bayesian update of the time t commitment prior) and the time t view of S_{t+2} . The weight on the former is non-negative but the weight λ_{t+1} on the latter could be negative. This is ruled out under Positive Prior-Bias but is compatible with Negative Prior-Bias.

These functional forms for q_{t+1} support our choice of terminology. When $q_{t+1} = p_{t+1}$, updating consists of applying Bayes' Rule to the commitment prior, which embodies "the correct" combination of prior beliefs and responsiveness to data. On the other hand, using the ex ante view expressed by $\sum_{s'_{t+1}} m_t(s'_{t+1}) p_{t+1}(\cdot | s'_{t+1})$ as the posterior would give all the weight to prior beliefs and none to data because the ex ante view does not depend on s_{t+1} . Thus an agent who updates according to the average scheme in (4.6) exhibits a positive bias to the prior if $\lambda_{t+1} > 0$ and a negative one if $\lambda_{t+1} < 0$.

Though q_{t+1} describes urges for making choices at $t + 1$, the agent balances it with the commitment view represented by p_{t+1} , as described in Section 2.3, and acts as though she forms the posterior one-step-ahead belief $m_{t+1} = \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}}$. The above noted bias of q_{t+1} extends to this mixture of p_{t+1} and q_{t+1} . For another angle on this interpretation, substitute for q_{t+1} from (4.6) and deduce that

$$m_{t+1} = \left(1 - \frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}}\right) p_{t+1} + \frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}} \left[\sum_{s'_{t+1}} m_t(s'_{t+1}) p_{t+1}(\cdot | s'_{t+1})\right]. \quad (4.7)$$

If $\lambda_{t+1} \geq 0$, one can think of the agent as overlooking the evidence represented by s_{t+1} with probability $\frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}}$, in which case she continues to use her time t beliefs about S_{t+2} .¹²

Note that (4.6) 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 and the process $(p_t, \alpha_t, \lambda_t)$.

¹²We considered naming the above axioms 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 (4.7) unless α_{t+1} and λ_{t+1} do not depend on s_{t+1} . See Section 5.1 for more on underreaction and overreaction.

Further content can be introduced into the model described in (4.6) 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.

4.2. Sample-Bias

In the last section, temptation and hence also the updating bias, depended on ex ante beliefs. Here we describe an alternative specialization of the general model in which temptation and the updating bias depend instead on sample frequencies.

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} . Let G lie in \mathcal{C} . Then $G(s_{t+2})$ is a subset of $C \times \mathcal{C}$ and so is the mixture $\int G(s'_{t+2}) d\Psi_{t+1}$. Consider the contingent menu in \mathcal{C}_{t+1} , denoted $\int G d\Psi_{t+1}$, that assigns $\int G(s'_{t+2}) d\Psi_{t+1}$ to every s_{t+2} . Then $(c_{t+1}, \int G d\Psi_{t+1})$ denotes the obvious singleton menu.

The axioms to follow parallel the trio of axioms stated in the last section. One difference is that the contingent menus F and G appearing in these axioms are assumed, for reasons given below, to lie in $\mathcal{C}_{+1}^0 \subset \mathcal{C}$. Thus F and G provide perfect commitment and are such that all relevant uncertainty is resolved by $t + 2$.

Axiom 15 (Sample-Bias). *For all s_{t+1} and c_{t+1} , for all F' in \mathcal{C} , and for all F and G in \mathcal{C}_{+1}^0 : if*

$$\begin{aligned} & (F'_{-s_{t+1}}, (c_{t+1}, F)) \succ_t (F'_{-s_{t+1}}, (c_{t+1}, G)) \text{ and} \\ & \left(F'_{-s_{t+1}}, (c_{t+1}, \int F d\Psi_{t+1}) \right) \sim_t \left(F'_{-s_{t+1}}, (c_{t+1}, \int G d\Psi_{t+1}) \right), \end{aligned} \quad (4.8)$$

then

$$(F'_{-s_{t+1}}, (c_{t+1}, F)) \sim_t (F'_{-s_{t+1}}, \{(c_{t+1}, F), (c_{t+1}, G)\}).$$

The next two axioms provide alternative strengthenings of Sample-Bias. Label by **Positive Sample-Bias** the axiom obtained if (4.8) is replaced by

$$\left(F'_{-s_{t+1}}, (c_{t+1}, \int F d\Psi_{t+1}) \right) \succeq_t \left(F'_{-s_{t+1}}, (c_{t+1}, \int G d\Psi_{t+1}) \right). \quad (4.9)$$

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

$$\left(F'_{-s_{t+1}}, (c_{t+1}, \int F d\Psi_{t+1}) \right) \succeq_t \left(F'_{-s_{t+1}}, (c_{t+1}, \int G d\Psi_{t+1}) \right). \quad (4.10)$$

Interpret Positive Sample-Bias; the other interpretations are similar. First, we interpret (4.9) 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 $\int F d\Psi_{t+1}$. But the agent is indifferent between F and its average because she satisfies Independence. Thus (4.9) implies that the average for F is better than that of G . Now 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 , 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 $\int F d\Psi_{t+1}$, relies on F lying in \mathcal{C}_{+1}^0 . 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}_{+1}^0 , however, information is irrelevant because all uncertainty is resolved once s_{t+2} is realized.

Corollary 4.2. *Let (\succeq_t) satisfy the axioms in Theorem 3.1. Then it satisfies Sample-Bias if and only if it admits representation as in the Theorem where in addition*

$$q_{t+1}(\cdot | s_{t+1}) = (1 - \lambda_{t+1}) p_{t+1}(\cdot | s_{t+1}) + \lambda_{t+1} \Psi_{t+1}(\cdot), \quad (4.11)$$

for some adapted process (λ_t) with $\lambda_{t+1} \leq 1$.¹³

Further, (\succeq_t) satisfies (i) Positive Sample-Bias or (ii) Negative Sample-Bias if and only if (4.11) is satisfied with respectively (i) $0 \leq \lambda_{t+1} \leq 1$ and (ii) $\lambda_{t+1} \leq 0$.

The implications of the functional form (4.11) are best seen through the implied adjustment rule for one-step-ahead beliefs, which has the form

¹³When $\lambda_{t+1} < 0$ in (4.11), q_{t+1} is well-defined as a probability measure only under special conditions; for example, it suffices that $\frac{-\lambda_{t+1}}{1-\lambda_{t+1}} \leq \min_{s_{t+2}} p_{t+1}(s_{t+2} | s_{t+1})$.

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

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 Bias,

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

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*.

Because she uses the empirical frequency measure to summarize past observations, the temptation facing an agent satisfying any of the models in the above corollary 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. For example, both the interpretations of the above axioms and the corollary remain intact if Ψ_{t+1} is 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 $w_{\tau, t+1} = \frac{1}{t+1}$ for all τ yields the earlier model. Thus the framework, including axiomatic foundations, permits a large variety of biases due to undue influence of the sample. For example, an agent who is influenced only by the most recent observation is captured by the law of motion

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

If $\lambda_{t+1} < 0$, the resulting model admits interpretation (in terms of sampling without replacement from changing urns) analogous to that offered by Rabin [19] for his model of the law of small numbers.

5. LEARNING ABOUT PARAMETERS

This section describes an example of our model in which the data generating process is unknown up to a parameter $\theta \in \Theta$. In the benchmark Bayesian model,

time t beliefs have the form

$$P_t(\cdot) = \int_{\Theta} \otimes_{t+1}^T \ell(\cdot | \theta) d\mu_t, \quad (5.1)$$

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 (5.1) that accommodates non-Bayesian updating.

Our specialization of the model (2.5)-(2.7) to accommodate parameters is defined by 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

$$\mu_{t+1} = \frac{BU(\mu_t; s_{t+1}) + \alpha_{t+1} \nu_{t+1}}{1 + \alpha_{t+1}}, \quad (5.2)$$

where $BU(\mu_t; s_{t+1})(\cdot)$ is the Bayesian update of μ_t .¹⁴ This equation constitutes the *law of motion* for beliefs about parameters. Finally, define (p_{t+1}, q_{t+1}) by

$$p_{t+1}(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d(BU(\mu_t; s_{t+1})) \quad \text{and} \quad (5.3)$$

$$q_{t+1}(\cdot) = \int_{\Theta} \ell(\cdot | \theta) d\nu_{t+1}. \quad (5.4)$$

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

Notice that

$$m_{t+1}(\cdot) = \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}} = \int_{\Theta} \ell(\cdot | \theta) d\mu_{t+1}. \quad (5.5)$$

In light of the discussion following (2.10), both the ranking of one-step-ahead bets and the choice out of menus at $t + 1$ are based on the beliefs about parameters

¹⁴ $d[BU(\mu_t; s_{t+1})](\theta) = \frac{\ell(s_{t+1}|\theta) d\mu_t(\theta)}{\int \ell(s_{t+1}|\theta') d\mu_t(\theta')}$.

represented by μ_{t+1} . If $\alpha_{t+1} \equiv 0$, then (μ_t) is the process of Bayesian posteriors and the above collapses to the exchangeable model (5.1). More generally, differences from the Bayesian model depend on (ν_t) , examples of which are given next.¹⁵

5.1. Prior-Bias with Parameters

Consider first the case where

$$\nu_{t+1} = (1 - \lambda_{t+1}) BU(\mu_t; s_{t+1}) + \lambda_{t+1}\mu_t, \quad (5.6)$$

where $\lambda_{t+1} \leq 1$. This is readily seen to imply (4.6) 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

$$\mu_{t+1} = \left(1 - \frac{\alpha_{t+1}\lambda_{t+1}}{1+\alpha_{t+1}}\right) BU(\mu_t; s_{t+1}) + \frac{\alpha_{t+1}\lambda_{t+1}}{1+\alpha_{t+1}}\mu_t. \quad (5.7)$$

The latter equation reveals something of how the inferences of an agent with Prior-Bias differ from those of a Bayesian updater. Compute that (assuming $\alpha_{t+1} \neq 0$)

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} < \frac{\ell(s_{t+1}|\theta)}{\ell(s_{t+1}|\theta')} \frac{\mu_t(\theta)}{\mu_t(\theta')} \quad \text{iff} \quad \lambda_{t+1}\ell(s_{t+1}|\theta') < \lambda_{t+1}\ell(s_{t+1}|\theta). \quad (5.8)$$

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 a Positive Prior-Bias (positive λ 's), then repeated application of (5.8) 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.

Turn to the question of what is learned in the long run. Learning may either signify learning the true parameter or learning to forecast future outcomes.¹⁶ The latter kind of learning is more relevant to choice behavior and thus is our focus. Suppose that $\theta^* \in \Theta$ is the true parameter and thus that the i.i.d. measure

¹⁵One 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.

¹⁶See [14] for the distinction between these two kinds of learning.

$P^* = \otimes_{t=1}^{\infty} \ell(\cdot | \theta^*)$ is the probability law describing the process (s_t) . Say that forecasts are *eventually correct on a path* s_1^∞ if, along that path,

$$m_t(\cdot) \longrightarrow \ell(\cdot | \theta^*) \quad \text{as } t \longrightarrow \infty.$$

Rewrite the law of motion for posteriors (5.7) in the form

$$\mu_{t+1} = (1 - \gamma_{t+1}) BU(\mu_t; s_{t+1}) + \gamma_{t+1} \mu_t, \quad (5.9)$$

where $\gamma_{t+1} = \frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}} \leq 1$. In general, γ_{t+1} is \mathcal{S}_{t+1} -measurable (γ_{t+1} may depend on the entire history s_1^{t+1} , including s_{t+1}), but we will be interested also in the special case where γ_{t+1} is \mathcal{S}_t -measurable. In that case, (5.9) can be interpreted not only in terms of Positive and Negative Prior-Bias as above, but also in terms of underreaction and overreaction to data. For example, let $\gamma_{t+1} \geq 0$ (corresponding to $\lambda_{t+1} \geq 0$). Then μ_{t+1} is a mixture, with weights that are independent of s_{t+1} , of two terms: (i) the Bayesian update $BU(\mu_t; s_{t+1})$, which incorporates the ‘correct’ response to s_{t+1} , and (ii) the prior μ_t , which does not respond to s_{t+1} at all. In a natural sense, therefore, an agent with $\gamma_{t+1} \geq 0$ *underreacts* to data. Similarly, if $\gamma_{t+1} \leq 0$, then $BU(\mu_t; s_{t+1})$ is a mixture of μ_{t+1} and μ_t , which suggests that μ_{t+1} reflects *overreaction*. Clearly, if $\gamma_{t+1} = 0$ then the model reduces to the Bayesian updating rule.

Theorem 5.1. *Let Θ be finite and $\mu_0(\theta^*) > 0$.*

(a) *Suppose that γ_{t+1} is \mathcal{S}_t -measurable and that $\gamma_{t+1} \geq 0$. Then forecasts are eventually correct P^* – a.s.*

(b) *Suppose that γ_{t+1} is \mathcal{S}_t -measurable and that $\gamma_{t+1} \leq 1 - \epsilon$ for some $\epsilon > 0$. Then forecasts are eventually correct with P^* -strictly positive probability.*

(c) *If one drops either of the assumptions in (a), then there exist (S, Θ, ℓ, μ_0) and $\theta \neq \theta^*$ such that*

$$m_t(\cdot) \longrightarrow \ell(\cdot | \theta) \quad \text{as } t \longrightarrow \infty,$$

with P^ -strictly positive probability.*

Assume that before any data are observed the prior belief puts positive weight on the true parameter, that is, assume that $\mu_0(\theta^*) > 0$. Then multiple repetition of Bayes’ Rule leads to near correct forecasts. This result is central in the Bayesian literature because it shows that the mere repetition of Bayes’ Rule eventually transforms the historical record into a near perfect guide for the future. Part (a)

of the theorem generalizes the Bayesian result to the case of underreaction. This result shows that, if repeated sufficiently many times, all non-Bayesian updating rules in (5.9) with the additional proviso of a Positive Prior-Bias and the indicated added measurability assumption, eventually produce good forecasting. Hence, in the case of underreaction, agent's forecasts converge to rational expectations although the available information is not processed according to Bayesian laws of probability.

Part (b) shows that, with positive probability, forecasts are eventually correct provided that the Bayesian term on the right side of (5.9) receives weight that is bounded away from zero. This applies in the case of Negative Prior-Bias, corresponding to overreaction. In fact, the results holds even if the forecaster sometimes overreacts and sometimes underreacts to new information. However, part (c) shows that convergence to wrong forecasts may occur in the absence of either of the assumptions in (a). This is demonstrated by two examples. In the first example the weight γ_{t+1} is constant, but sufficiently negative, corresponding to a forecaster that sufficiently overreacts to new information. In the second example, the weight γ_{t+1} is positive corresponding to underreaction, but γ_{t+1} depends on the current signal and, therefore, γ_{t+1} is only \mathcal{S}_{t+1} -measurable. In both examples, forecasts may eventually converge to an incorrect limit. Moreover, wrong forecasts in the limit are at least as likely to occur as are correct forecasts.

The proof of Theorem 5.1 builds on classic arguments of the Bayesian literature. Consider the probability measure μ_t on the parameter space and let the random variable μ_t^* be the probability that μ_t assigns to the true parameter. It follows that the expected value (according to the true data generating process) of the Bayesian update of μ_t^* (given new information) is greater than μ_t^* itself. Hence, in the Bayesian case, the weight given to the true parameter tends to grow as new information is observed. This submartingale property ensures that Bayesian forecasts must converge to some value and cannot remain in endless random fluctuations. The submartingale property follows because under the Bayesian paradigm future changes in beliefs that can be predicted are incorporated in current beliefs. It is immediate from the linear structure in (5.9) that this basic submartingale property still holds in our model as long as the weight γ_{t+1} depends upon the history only up to period t . Hence, with this measurability assumption, forecasts in our model must also converge and, as in the Bayesian case, cannot remain in endless random fluctuations.¹⁷ In addition, convergence to the truth holds in

¹⁷We conjecture that beliefs μ_t may not converge in some examples when the weight γ_{t+1} is \mathcal{S}_{t+1} -measurable. In our example, it does converge, but to an incorrect limit.

both the Bayesian paradigm and in the case of underreaction. However, given sufficiently strong overreaction, it is possible that forecasts will settle on an incorrect limit. This follows because the positive drift of the above mentioned submartingale property on μ_t^* may be compensated by sufficiently strong volatility which permits that, with positive probability, μ_t^* converges to zero.

5.2. Sample-Bias with Parameters

Sample-Bias can also be modeled when learning about parameters is taking place. 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

$$\Psi_{t+1} = \int \ell(\cdot | \theta) d\psi_{t+1}(\theta). \quad (5.10)$$

Let μ_0 be given and define μ_{t+1} and ν_{t+1} inductively for $t \geq 0$ by (5.2) and

$$\nu_{t+1} = (1 - \lambda_{t+1}) BU(\mu_t, s_{t+1}) + \lambda_{t+1}\psi_{t+1}, \quad (5.11)$$

for $\lambda_{t+1} \leq 1$. Then one obtains a special case of the Sample-Bias model of Corollary 4.2; the bias is positive or negative according to the sign of the λ 's. The implied law of motion for posteriors is

$$\mu_{t+1} = \left(1 - \frac{\alpha_{t+1}\lambda_{t+1}}{1+\alpha_{t+1}}\right) BU(\mu_t; s_{t+1}) + \frac{\alpha_{t+1}\lambda_{t+1}}{1+\alpha_{t+1}}\psi_{t+1}. \quad (5.12)$$

To illustrate, suppose that $S = \{s^1, \dots, s^K\}$ and that $\ell(s^k | \theta) = \theta_k$ for each $\theta = (\theta_1, \dots, \theta_K)$ in Θ , the interior of the K -simplex. Then one can ensure (5.10) 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

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

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,

$$\frac{d\psi_t(\theta)}{\prod_{k=1}^K d\theta_k} \propto \prod_{k=1}^K \theta_k^{n_t(s^k)-1}. \quad (5.13)$$

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 (5.10).

Finally, compute from (5.12) and (5.13) that (assuming $\alpha_{t+1} \neq 0$)

$$\frac{\mu_{t+1}(\theta)}{\mu_{t+1}(\theta')} > \frac{\ell(s_{t+1}|\theta)}{\ell(s_{t+1}|\theta')} \frac{\mu_t(\theta)}{\mu_t(\theta')} \quad \text{iff} \quad \lambda_{t+1} \frac{\psi_t(\theta)}{\psi_t(\theta')} > \lambda_{t+1} \frac{\mu_t(\theta)}{\mu_t(\theta')}. \quad (5.14)$$

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 (5.13) and (5.14) 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 if the common initial prior μ_0 is uniform.¹⁸ 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.

We conclude with a result regarding learning in the long run. In order to avoid technical issues arising from Θ being a continuum as in the Dirichlet-based model, we consider the following variation: as before $S = \{s^1, \dots, s^K\}$ and $\ell(s^k | \theta) = \theta_k$

¹⁸More generally, the latter two conditions can be replaced by $\frac{\theta'(1-\theta')}{\theta(1-\theta)} > \frac{\mu_0(\theta)}{\mu_0(\theta')}$.

for each k and θ . But now take Θ to be the set of points $\theta = (\theta_1, \dots, \theta_K)$ in the interior of the K -simplex having rational co-ordinates. Define

$$\psi_{t+1}(\theta) = \begin{cases} 1 & \text{if the empirical frequency of } s^k \text{ is } \theta_k, 1 \leq k \leq K, \\ 0 & \text{otherwise.} \end{cases}$$

Then (5.10) is evident.¹⁹ The law of motion can be written in the form

$$\mu_{t+1} = (1 - \gamma_{t+1}) BU(\mu_t; s_{t+1}) + \gamma_{t+1} \psi_{t+1}, \quad (5.15)$$

where $\gamma_{t+1} = \frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}} \leq 1$.

We have the following partial counterpart of Theorem 5.1.

Theorem 5.2. *Let S , (Θ, ℓ) and (ψ_t) be as just defined and suppose that posteriors (μ_t) evolve according to (5.15), where γ_{t+1} is \mathcal{S}_t -measurable and $0 < \underline{\gamma} \leq \gamma_{t+1} \leq 1$. Then forecasts are eventually correct P^* - a.s.*

The positive lower bound $\underline{\gamma}$ excludes the Bayesian case. The result does hold in the Bayesian case $\gamma_{t+1} = 0$. However, unlike the proof of Theorem 5.1, the proof of Theorem 5.2 is in some ways significantly different from the proof in the Bayesian case. We suspect that the differences in the approach make the lower bound assumption technically convenient but ultimately disposable. We also conjecture (but cannot yet prove) that just as in part (c) of Theorem 5.1, convergence to the truth fails in general if γ_{t+1} is only \mathcal{S}_{t+1} -measurable. The other case treated in the earlier theorem - γ_{t+1} is \mathcal{S}_t -measurable but possibly negative - (which in the context of that model corresponded to overreaction) is not relevant here because these conditions violate the requirement that each ν_{t+1} in (5.11) be a probability measure and hence non-negatively valued.

6. CONCLUDING REMARKS

We conclude by addressing questions that may have occurred to some readers and by pointing to some possible extensions.

¹⁹If Θ were taken to be finite, then one could not assure (5.10) without admitting signed measures for ψ_{t+1} and hence also for μ_{t+1} . Bayesian updating is not well-defined for signed measures and even if that problem were overcome, the interpretation of such a model is not clear.

Can non-Bayesian updaters survive? A traditional objection to the use of non-Bayesian updating rules is that they would render an agent vulnerable to a Dutch Book or money pump. Dutch Book arguments are typically loose but they center on showing that non-Bayesian updating leads to dynamic inconsistency which in turn permits the agent to be exploited. In contrast, we have shown in our explicitly dynamic formal setup that our agent is guided through time by a single (complete and transitive) preference order. Thus he is dynamically consistent and immune to Dutch Books.

Another objection comes from the efficient markets hypothesis. This objection is based on the assumption that agents who do not process information according to Bayes' Rule will not accumulate sufficient wealth to have a significant impact on asset prices. Sandroni [20] explores this issue but the preferences of agents in his model are different from those described here. Whether or not Bayesians accumulate more wealth than our non-Bayesian updaters is an open question. We would add that the significance of non-Bayesian agents may not be dismissed even if they accumulate less wealth than Bayesians. First, some wealth may belong to all types of agents if there is continuous entry into the market. Second, trade restrictions, such as the impossibility of selling future wages, imply that all agents will keep their wages and so will, at least, influence economic quantities and the prices of goods.²⁰

A matter of framing: The agent in our model evaluates prospects by backward induction, thus taking into account future updating behavior. It merits emphasis that she does this even when evaluating commitment prospects ($F \in \mathcal{C}^0$). One might view this feature of the model as problematic because commitment prospects do not involve interim choice and hence do not call *explicitly* for future updating. However, contingent menus are dynamic objects and thus backward induction reasoning is intuitive in our view. It is true that there is a natural identification on formal grounds between contingent menus in \mathcal{C}^0 and *acts* in \mathcal{A} , the set of all (suitably measurable) mappings $f : S^\infty \rightarrow C^\infty$. But we distinguish between contingent menus and acts on cognitive grounds - contingent menus are dynamic objects while acts are static and thus it is plausible that the agent evaluate them differently. The implicit distinction between \mathcal{C}^0 and \mathcal{A} is the way in which our model acknowledges that *framing* matters for choice. To make this

²⁰In addition, the existence of a direct link between wealth and influence on asset prices have been questioned by Kogan et al [11].

distinction explicit and formal, one would need to consider not only preferences \succeq_t on \mathcal{C} but also preferences \succeq_t^A on \mathcal{A} . Such an extension is a subject for future research.

To illustrate how framing might matter, consider uncertainty described by 2 tosses of a coin, with the tosses coming at $t = 1, 2$. The agent is told the following at $t = 0$: “Your consumption in periods 0 and 1 is fixed; consumption in other periods is constant across time but depends on which of the prospects A and B you choose. If H_1 (the first toss yields Heads), then: A pays high consumption if H_2 and low otherwise, while B pays low consumption if H_2 and high otherwise. If T_1 on the other hand, then A pays high only in T_2 and B only in H_2 . Now choose between between A and B .”

One modeling hypothesis is that the agent perceives the prospects as acts (elements of \mathcal{A}), where A is a bet on the two tosses producing identical outcomes and B is a bet on the two outcomes being different. Then her evaluation would depend only on her initial prior on $S_1 \times S_2$, the obvious i.i.d. prior if the coin is thought to be unbiased, and updating would be irrelevant. On the other hand, the above description frames A and B as dynamic prospects (elements of \mathcal{C}^0) - it calls attention to future conditional positions and may bring to mind consideration of future updating and backward induction. Our agent sees the prospects in this way. Thus the measure P_0 from (2.10) that she uses to evaluate commitment prospects at time 0 reflects not only prior beliefs about the coin but also future updating biases.

Other extensions: Our main contribution is to provide a choice-theoretic model of updating. An important feature of the model is its richness - it can accommodate a range of updating biases. We have illustrated this to a degree via the (axiomatic) specializations called Prior-Bias and Sample-Bias. However, much more might be done in this vein. For example, we characterized two alternative specializations of Sample-Bias that correspond roughly to the hot-hand fallacy (Positive Sample-Bias) and the gambler’s fallacy (Negative Sample-Bias) respectively. However, while in each case the agent is assumed to suffer from the indicated fallacy at all times and histories, 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: in light of Corollary 4.2, one must ‘only’ explain how the weights λ_{t+1} vary with history. Similarly with regard to further specializations of Prior-Bias.

A. APPENDIX: Contingent Menus

Define the following spaces:

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

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

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 $\mathcal{C} \subset \Pi_1^\infty D_t$ 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 $\Pi_1^\infty D_t$ into D_t . Then $\pi_t(\mathcal{C}) \subset \mathcal{C}$ and

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

- (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),

$$(\lambda \circ F' + (1 - \lambda) \circ F)(s) =$$

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

Part (i) asserts that the topological structure of C is inherited by \mathcal{C} . Part (ii) is the homeomorphism (2.2) that was used heavily in the text.

We noted 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 3. Roughly it shows that ‘ \circ ’, which is the natural mixing operation induced by $\mathcal{C} \subset \Pi_1^\infty D_t$, 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 $\lambda \circ F' + (1 - \lambda) \circ F$.

Parts (i)-(iv) of the theorem are closely related to several results in the literature dealing with hierarchies of topological spaces and problems of infinite regress. For example, Mertens and Zamir [16] and Brandeburger and Dekel [2] study hierarchies of probability measures, Epstein and Wang [7] study hierarchies of preferences and GP [10] establish a recursive domain suitable for their infinite horizon model. The technical details are now well understood and thus we omit a formal proof. Note, moreover, that the result for singleton S , when we are simply dealing with hierarchies of closed subsets, is a corollary of [7, Theorem 6.1]. See also [6, 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.

Finally, define the spaces $\mathcal{C}^r \subset \mathcal{C}_{+1}^0 \subset \mathcal{C}^0 \subset \mathcal{C}$. First, \mathcal{C}^0 is the unique subspace of \mathcal{C} satisfying: $\mathcal{C}^0 \underset{\text{homeo}}{\approx} (C \times \mathcal{C}^0)^S$ under the homeomorphism in the theorem. (Details are as in [7, Theorem 6.1(a)].) Take $\mathcal{C}_{+1}^0 = D_1$. Finally, let $\mathcal{C}^r = \mathcal{C}^\infty$, which can be identified with a subset of \mathcal{C} ; $\mathcal{C}^r \underset{\text{homeo}}{\approx} C \times \mathcal{C}^r$ under the induced homeomorphism.

B. APPENDIX: Proof of Main Theorem

Necessity: 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

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

Then Γ is a contraction under assumption (3.13) and thus it has a unique fixed point (U_t) . Define

$$\mathcal{U}_t(F_t) = \int_{S_{t+1}} U_{t+1}(F_t(s_{t+1}), s_{t+1}) dm_t(s_{t+1}), \quad F_t \in \mathcal{C}_t.$$

It remains to verify Axioms 1-13.

Order and Continuity are by construction. Nondegeneracy follows from non-constancy of u . For Independence, let X^L consist of all (U_t) in X such that each $U_t(\cdot, s_1^t)$ is mixture linear. Then X^L is closed in X and Γ maps X^L into itself. Thus the unique fixed point noted above is in X^L . Independence follows.

Some abbreviations are used in the sequel. Define the conditional order $\succeq_t|_{s_{t+1}}$ on $\mathcal{M}(\mathcal{C}_{t+1} \times \mathcal{C}_{t+1})$ by $M' \succeq_t|_{s_{t+1}} M$ iff $\exists F$ such that $(F_{-s_{t+1}}, M') \succeq_t (F_{-s_{t+1}}, M)$. Define $\succeq_t|_{s_{t+1}, s_{t+2}}$ on closed subsets of $\mathcal{M}(\mathcal{C}_{t+2} \times \mathcal{C}_{t+2})$ by: $L' \succeq_t|_{s_{t+1}, s_{t+2}} L$ iff there exist F, c and $M \subset \mathcal{C}$ such that

$$(F_{-s_{t+1}}, (c, L's_{t+2}M)) \succeq_t (F_{-s_{t+1}}, (c, Ls_{t+2}M)).$$

Several of the axioms can be restated using these derived orders.

Absolute Continuity: Hypothesis (3.5) $\iff m_t(s_{t+1})p_{t+1}(s_{t+2}) = 0 \iff m_t(s_{t+1})m_{t+1}(s_{t+2}) = 0$, (because $q_{t+1} \ll p_{t+1}$), $\iff (s_{t+1}, s_{t+2}) \succeq_t$ -null \implies (3.6).

Verify that $\succeq_t|_{s_{t+1}}$ is represented by $m_t(s_{t+1})U_{t+1}(\cdot, s_{t+1})$. It follows that \succeq_t satisfies Set-Betweenness. Verify also that $\succeq_t|_{s_{t+1}, s_{t+2}}$ is represented by

$$L \longmapsto m_t(s_{t+1})p_{t+1}(s_{t+2}) \max_{\ell \in L} U_{t+2}(\ell, s_{t+2}). \quad (\text{B.1})$$

State Independence: Follows from (B.1) and the fact that $U_{t+2}(\ell, s_{t+2})$ does not depend on the state if $\ell \in \mathcal{C} \times \mathcal{C}^r$.

RSR: Follows from (B.1).

$$\text{CSR: } (K', F') \succeq_{t|s_{t+1}} (K, F') \iff m_t(s_{t+1}) \max_{c \in K'} u(c) \geq m_t(s_{t+1}) \max_{c \in K} u(c) \implies \\ m_t(s_{t+1}) \max_{c \in K'} u(c) = m_t(s_{t+1}) \max_{c \in K' \cup K} u(c) \implies (K', F') \sim_{t|s_{t+1}} (K' \cup K, F').$$

Risk Preference: Obvious.

Invariant Discounting: Under the indicated temptation hypothesis,

$$U_{t+1} \left(\left\{ (c, F'), (d, \vec{d}) \right\}, s_{t+1} \right) = (1 + \alpha_{t+1}) u(c) + \delta \int U_{t+2}(F'(s_{t+2})) dm_{t+1} - \\ \alpha_{t+1} U_{t+1} \left((d, \vec{d}), s_{t+1} \right), \text{ and similarly for } (e, \vec{e}).$$

RRC: Let (s_{t+1}, s_{t+2}) be \succeq_t -non-null. Then $\succeq_{t|s_{t+1}, s_{t+2}}$ is represented by $L \mapsto \max_{\ell \in L} U_{t+2}(\ell, s_{t+2})$ and $\succeq_{t+1|s_{t+2}}$ is represented by $U_{t+2}(\cdot, s_{t+2})$. Thus, for $\ell', \ell \subset \mathcal{C}_{t+2} \times \mathcal{C}_{t+2}$, $\{\ell'\} \succeq_{t|s_{t+1}, s_{t+2}} \{\ell\}$ iff $U_{t+2}(\ell', s_{t+2}) \geq U_{t+2}(\ell, s_{t+2})$ iff $\ell' \succeq_{t+1|s_{t+2}} \ell$.

Bias Persistence: If (3.10), then

$$\int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) dp_{t+1} > \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dp_{t+1}$$

and

$$\int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) dq_{t+1} \geq \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dq_{t+1}.$$

Therefore,

$$\int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1} q_{t+1}] \geq \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1} q_{t+1}] \\ \implies \int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) dm_{t+1} \geq \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dm_{t+1} \\ \implies \mathcal{U}_{t+1}(F') \geq \mathcal{U}_{t+1}(G).$$

For the other case, the self-control hypothesis (3.11) implies that

$$U_{t+1}((c, \{G, H\}), s_{t+1}) = u(c) + \delta \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dp_{t+1} + \\ \alpha_{t+1} \delta \left\{ \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dq_{t+1} - \int_{S_{t+2}} U_{t+2}(H(s_{t+2}), s_{t+2}) dq_{t+1} \right\},$$

and similarly for F' in place of G . Thus

$$\{(c, F'), (c, H)\} \succeq_{t|s_{t+1}} \{(c, G), (c, H)\} \implies$$

$$\begin{aligned}
\int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1} q_{t+1}] &\geq \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1} q_{t+1}] \\
\implies \int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) dm_{t+1} &\geq \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dm_{t+1} \\
&\implies \mathcal{U}_{t+1}(F') \geq \mathcal{U}_{t+1}(G).
\end{aligned}$$

Sufficiency: The argument is divided into a series of steps. Steps 1 and 2 are similar to arguments in [5] for the 3-period case and thus some details are omitted.

Step 1: $\mathcal{U}_t(F) = \Sigma_{s_{t+1}} U_{t+1}^*(F(s_{t+1}), s_{t+1})$, where $U_{t+1}^*(\cdot, s_{t+1})$ is linear on $\mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1})$. This follows from Order, Continuity and Independence for \succeq_t on \mathcal{C}_t .

It follows that $\succeq_t|_{s_{t+1}}$ is complete and transitive and represented by $U_{t+1}^*(\cdot, s_{t+1})$. The latter is not constant if s_{t+1} is \succeq_t -non-null. By Nondegeneracy, there is at least one non-null state for each \succeq_t .

Step 2: $\succeq_t|_{s_{t+1}}$ satisfies GP axioms suitably translated to $\mathcal{M}(C_{t+1} \times \mathcal{C}_{t+1})$. Thus

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

for some $U_{t+1}^{GP}(\cdot, s_{t+1})$ and $V_{t+1}^{GP}(\cdot, s_{t+1})$, linear functions on $C_{t+1} \times \mathcal{C}_{t+1}$. It follows that these functions are additive across states in S_{t+2} in the sense that

$$U_{t+1}^{GP}(c, F, s_{t+1}) = \Sigma_{s_{t+2}} u_t(c, F(s_{t+2}), s_{t+1}, s_{t+2}) \quad (\text{B.2})$$

and

$$V_{t+1}^{GP}(c, F, s_{t+1}) = \Sigma_{s_{t+2}} v_t(c, F(s_{t+2}), s_{t+1}, s_{t+2}), \quad (\text{B.3})$$

where each $u_t(\cdot, s_{t+1}, s_{t+2})$ and $v_t(\cdot, s_{t+1}, s_{t+2})$ is linear and continuous on $C_{t+1} \times \mathcal{M}_{t+2}$. The subscript t indicates that these functions may depend also on the history s_1^t underlying \succeq_t .

Step 3: Derive further structure for u_t and v_t .

Define $\succeq_t|_{s_{t+1}, s_{t+2}}$ as in the proof of necessity. From (B.2)-(B.3), deduce that $\succeq_t|_{s_{t+1}, s_{t+2}}$ is represented numerically by $\tilde{U}_t(\cdot, s_{t+1}, s_{t+2})$, where

$$\tilde{U}_t(L, s_{t+1}, s_{t+2}) = \max_{\ell \in L} \{u_t(c_{t+1}, \ell, s_{t+1}, s_{t+2}) + v_t(c_{t+1}, \ell, s_{t+1}, s_{t+2})\}$$

$$- \max_{\ell' \in L} v_t(c_{t+1}, \ell', s_{t+1}, s_{t+2}), \text{ for } L \subset \mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2}).$$

We claim that

$$v_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2}) = a_t(c_{t+1}, s_{t+1}, s_{t+2}) u_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2}) + A_t(c_{t+1}, s_{t+1}, s_{t+2}), \quad (\text{B.4})$$

for some $a_t \geq 0$.

Case 1: Suppose that $u_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2})$ is nonconstant. By RSR, $\succeq_t|_{s_{t+1}, s_{t+2}}$ is strategically rational. Deduce (B.4) as in [9, p. 1414].

Case 2: Suppose that $u_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2})$ is constant. Then Absolute Continuity implies that (s_{t+1}, s_{t+2}) is \succeq_t -null. But then *any* specification for $v_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2})$ is consistent with a representation for \succeq_t . In particular, we can take (B.4) to be valid with $a_t(c_{t+1}, s_{t+1}, s_{t+2}) = 0$ and some $A_t(c_{t+1}, s_{t+1}, s_{t+2})$.

From (B.4), deduce that $\succeq_t|_{s_{t+1}, s_{t+2}}$ is represented by

$$L \mapsto \max_{\ell \in L} u_t(c_{t+1}, \ell, s_{t+1}, s_{t+2}).$$

Let (s_{t+1}, s_{t+2}) be \succeq_t -non-null. Then RRC implies that

$$\ell' \succeq_{t+1|s_{t+2}} \ell \iff \{\ell'\} \succeq_t|_{s_{t+1}, s_{t+2}} \{\ell\},$$

for all ℓ', ℓ in $\mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$. By Step 1, $\succeq_{t+1|s_{t+2}}$ is represented by $U_{t+2}^*(\cdot, s_{t+2})$. Conclude that the latter is ordinally equivalent to $u_t(c_{t+1}, \cdot, s_{t+1}, s_{t+2})$ on $\mathcal{M}(C_{t+2} \times \mathcal{C}_{t+2})$ for every c_{t+1} . Since both are linear functions, they must be cardinally equivalent, that is,

$$u_t(c, \cdot, s_{t+1}, s_{t+2}) = b_t(c, s_{t+1}, s_{t+2}) U_{t+2}^*(\cdot; s_{t+2}) + B_t(c, s_{t+1}, s_{t+2}), \quad (\text{B.5})$$

where $b_t \geq 0$. But both $u_t(\cdot, s_{t+1}, s_{t+2})$ and $U_{t+2}^*(\cdot; s_{t+2})$ are linear; and (s_{t+1}, s_{t+2}) non-null implies that $u_t(c, \cdot, s_{t+1}, s_{t+2})$ and $U_{t+2}^*(\cdot, s_{t+2})$ are nonconstant. It follows that $b_t(\cdot, s_{t+1}, s_{t+2}) > 0$. Further, $b_t(\cdot, s_{t+1}, s_{t+2})$ is constant and $B_t(\cdot, s_{t+1}, s_{t+2})$ is linear.

To support the two latter claims, adopt the obvious simplified notation and let

$$u(c, M) = b(c) U(M) + B(c),$$

where: $u(\cdot)$ and $U(\cdot)$ are linear and U is not constant. Evaluation at $\lambda(c, M) + (1 - \lambda)(c', M')$ yields

$$[\lambda B(c) + (1 - \lambda) B(c') - B(\lambda c + (1 - \lambda) c')] =$$

$$\lambda U(M) (b(\lambda c + (1 - \lambda) c') - b(c)) + \\ (1 - \lambda) U(M') (b(\lambda c + (1 - \lambda) c') - b(c')).$$

If U is not constant, then one can vary $U(M)$ and $U(M')$ independently over open intervals. But the LHS is independent of M and M' . Deduce that $b(c) = b(c') = b(\lambda c + (1 - \lambda) c')$ and that $\lambda B(c) + (1 - \lambda) B(c') - B(\lambda c + (1 - \lambda) c') = 0$.

Equation (B.5) is valid also if (s_{t+1}, s_{t+2}) is \succeq_t -null: Nullity implies that $u_t(c, \cdot, s_{t+1}, s_{t+2})$ is constant. Hence take $b_t(\cdot, s_{t+1}, s_{t+2}) = 0$ and $B_t(c, s_{t+1}, s_{t+2}) = u_t(c, M_{t+2}, s_{t+1}, s_{t+2})$ for any M_{t+2} .

Step 4: Derive more structure for (u_t, v_t) and some implications.

The argument used for (B.5) can be applied to (B.4) to show that $a_t(\cdot, s_{t+1}, s_{t+2})$ is constant and that $A_t(\cdot, s_{t+1}, s_{t+2})$ is linear. It follows that

$$\mathcal{U}_t(F_t) = \sum_{s_{t+1}} U_{t+1}^*(F_t(s_{t+1}), s_{t+1}), \text{ where } U_{t+1}^*(M_{t+1}, s_{t+1}) = \\ \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ \begin{array}{l} \sum_{s_{t+2}} b_t(s_{t+1}, s_{t+2}) U_{t+2}^*(F_{t+1}(s_{t+2}), s_{t+2}) + \sum_{s_{t+2}} B_t(c_{t+1}, s_{t+1}, s_{t+2}) \\ + \sum_{s_{t+2}} v_t(c_{t+1}, F_{t+1}(s_{t+2}), s_{t+1}, s_{t+2}) \end{array} \right\} \\ - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \sum_{s_{t+2}} v_t(c'_{t+1}, F'_{t+1}(s_{t+2}), s_{t+1}, s_{t+2}),$$

where $v_t(\cdot, s_{t+1}, s_{t+2}) = a_t b_t U_{t+2}^*(F_{t+1}, s_{t+2}) + A'_t(c_{t+1}, s_{t+1}, s_{t+2})$, and $A'_t(c_{t+1}, s_{t+1}, s_{t+2}) = a_t B_t(c_{t+1}, s_{t+1}, s_{t+2}) + A_t(c_{t+1}, s_{t+1}, s_{t+2})$. Define $B''_t(c_{t+1}, s_{t+1}) = \sum_{s_{t+2}} B_t(c_{t+1}, s_{t+1}, s_{t+2})$ and $A''_t(c_{t+1}, s_{t+1}) = \sum_{s_{t+2}} A'_t(c_{t+1}, s_{t+1}, s_{t+2})$ and rewrite in the form:

$$U_{t+1}^*(M_{t+1}, s_{t+1}) = \\ \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ \begin{array}{l} \sum_{s_{t+2}} b_t(s_{t+1}, s_{t+2}) U_{t+2}^*(F_{t+1}(s_{t+2}), s_{t+2}) + B''_t(c_{t+1}, s_{t+1}) \\ + \sum_{s_{t+2}} a_t b_t U_{t+2}^*(F_{t+1}(s_{t+2}), s_{t+2}) + A''_t(c_{t+1}, s_{t+1}) \end{array} \right\} \\ - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \left[\sum_{s_{t+2}} a_t b_t U_{t+2}^*(F'_{t+1}(s_{t+2}), s_{t+2}) + A''_t(c'_{t+1}, s_{t+1}) \right]. \quad (\text{B.6})$$

The ranking of consumption menus $K \subset C$ induced by $\succeq_{t|s_{t+1}}$ is represented by

$$\max_{c_{t+1} \in K} \{B''_t(c_{t+1}, s_{t+1}) + A''_t(c_{t+1}, s_{t+1})\} - \max_{c'_{t+1} \in K} A''_t(c'_{t+1}, s_{t+1}).$$

By Consumption Strategic Rationality, this ranking is strategically rational. By Risk Preference, $B''_t(\cdot, s_{t+1})$ is not constant if s_{t+1} is \succeq_t -non-null. Then, as in [9, p. 1414],

$$A''_t(\cdot, s_{t+1}) = a'_t(s_{t+1}) B''_t(\cdot, s_{t+1}) + \bar{a}'_t(s_{t+1}),$$

where $a'_t(s_{t+1}) \geq 0$. Deduce that the restriction of $\succeq_{t|s_{t+1}}$ to consumption menus is represented by

$$K \longmapsto \max_{c_{t+1} \in K} B''_t(c_{t+1}).$$

By Risk Preference, this ranking is independent of t and the \succeq_t -non-null state s_{t+1} . It follows that, for any \succeq_t -non-null s_{t+1} ,

$$B''_t(\cdot, s_{t+1}) = a''_t(s_{t+1}) B^*(\cdot) + b''_t(s_{t+1}).$$

The preceding extends also to \succeq_t -null states s_{t+1} , in which case $U_{t+1}^*(\cdot, s_{t+1})$ is constant: take $a'_t(s_{t+1}) = a''_t(s_{t+1}) = 0$.

Finally, it follows from (B.6), after dropping irrelevant additive terms, that

$$\begin{aligned} & U_{t+1}^*(M_{t+1}, s_{t+1}) = \\ & \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ \begin{array}{l} (1 + a'_t(s_{t+1})) a''_t(s_{t+1}) B^*(c_{t+1}) + \\ \Sigma_{s_{t+2}} (1 + a_t(s_{t+1}, s_{t+2})) b_t(s_{t+1}, s_{t+2}) U_{t+2}^*(F_{t+1}(s_{t+2}), s_{t+2}) \end{array} \right\} \\ & - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \left[a'_t(s_{t+1}) a''_t(s_{t+1}) B^*(c'_{t+1}) + \Sigma_{s_{t+2}} a_t b_t U_{t+2}^*(F'_{t+1}(s_{t+2}), s_{t+2}) \right]. \end{aligned} \tag{B.7}$$

Step 5: Deliver probabilities for all t in spite of the pervasive state dependence. The idea is to first restrict attention to subclasses of contingent menus that mimic the objects of choice in a corresponding finite T -horizon setting. Here we can argue by backward induction - roughly speaking, state independence at terminal T provides a scale that permits defining meaningful probabilities at all prior times. Then we use Continuity to extend these probabilities to all of $\Pi_1^\infty S_t$.

Step 5(i): Fix $3 \leq T < \infty$ and define p_0 and $(p_t, q_t, \alpha_t)_1^{T-1}$ that represent preference in a sense to be described. Recall $D_T \subset \mathcal{C}$ from the previous appendix, the subdomain of contingent menus that imply commitment and no uncertainty (only risk) for periods $T+1$ and on. Until further notice, consider preferences restricted to D_T rather than defined on all of \mathcal{C} .

Argue by backward induction on t . As the induction hypothesis, suppose that for all τ , $t+1 \leq \tau \leq T-2$, \succeq_τ is represented by $\mathcal{U}_\tau(F_\tau) = \Sigma_{s_{\tau+1}} U_{\tau+1}^*(F_\tau(s_{\tau+1}), s_{\tau+1})$, where

$$U_{\tau+1}^*(F_\tau(s_{\tau+1}), s_{\tau+1}) = m'_\tau(s_{\tau+1}) U_{\tau+1}(F_\tau(s_{\tau+1}), s_{\tau+1}),$$

m'_τ is a measure on $S_{\tau+1}$, and where, for every \succeq_τ -non-null $s_{\tau+1}$, $U_{\tau+1}(\cdot, s_{\tau+1})$ has the form in (2.7). The only relation imposed on the relevant measures is that

$$q_{\tau+1} \ll p_{\tau+1} \quad \text{and} \quad p_{\tau+1} \ll m'_{\tau+1}. \quad (\text{B.8})$$

Thus the induction hypothesis amounts to the representation (2.6)-(2.7), where (B.8), but not (2.5), is assumed. We prove a corresponding representation for $\mathcal{U}_t(\cdot)$.

By definition, if $s_{\tau+2}$ is $\succeq_{\tau+1}$ -null, then the specification of $U_{\tau+2}(\cdot, s_{\tau+2})$ is of no consequence for $\mathcal{U}_{\tau+1}(\cdot)$ and hence also $\succeq_{\tau+1}$. In fact, it is of no consequence also for $\mathcal{U}_\tau(\cdot)$ or indeed for any $\mathcal{U}_k(\cdot)$: $s_{\tau+2}$ is $\succeq_{\tau+1}$ -null iff $m'_{\tau+1}(s_{\tau+2}) = 0$ and the latter implies, by (B.8), that $p_{\tau+1}(s_{\tau+2}) + \alpha_{\tau+1}q_{\tau+1}(s_{\tau+2}) = 0$, which means that $U_{\tau+2}(F_{\tau+1}(s_{\tau+2}), s_{\tau+2})$ enters into the appropriate version of (2.7) multiplied by zero. It follows that if $s_{\tau+2}$ is $\succeq_{\tau+1}$ -null, then the specification of $U_{\tau+2}(\cdot, s_{\tau+2})$ is of no consequence for the representation of \succeq_τ or indeed of any \succeq_k . Consequently, below we can restrict attention to specifying utilities at non-null states, secure in the knowledge that utilities at null states can be specified freely to satisfy the desired recursive equation (2.7).

From Steps 1 and 4,

$$\mathcal{U}_t(F_t) = \sum_{s_{t+1}} U_{t+1}^*(F_t(s_{t+1}), s_{t+1}),$$

where $U_{t+1}^*(M_{t+1}, s_{t+1})$ is given by (B.7). Define

$$b'_t(s_{t+1}) \equiv \sum_{s_{t+2}} b_t(s_{t+1}, s_{t+2}) m'_{t+1}(s_{t+2}).$$

Then $b'_t(s_{t+1}) = 0 \implies b_t(s_{t+1}, s_{t+2}) m'_{t+1}(s_{t+2}) = 0$ for all s_{t+2} . But then (B.7) and Risk Preference imply that s_{t+1} is \succeq_t -null.

For s_{t+1} that is \succeq_t -non-null, the preceding implies that

$$b'_t(s_{t+1}) > 0.$$

Define

$$\begin{aligned} p_{t+1}(s_{t+2}) &= \frac{b_t(s_{t+1}, s_{t+2}) m'_{t+1}(s_{t+2})}{b'_t(s_{t+1})}, \\ \alpha_{t+1} &= \frac{\sum_{s'_{t+2}} a_t(s_{t+1}, s'_{t+2}) b_t(s_{t+1}, s'_{t+2}) m'_{t+1}(s'_{t+2})}{b'_t(s_{t+1})} \geq 0, \\ q_{t+1}(s_{t+2}) &= \begin{cases} \frac{a_t(s_{t+1}, s_{t+2}) b_t(s_{t+1}, s_{t+2}) m'_{t+1}(s_{t+2})}{b'_t(s_{t+1}) \alpha_{t+1}} & \text{if } \alpha_{t+1} > 0 \\ p_{t+1}(s_{t+2}) & \text{otherwise;} \end{cases} \end{aligned}$$

note that $q_{t+1} \ll p_{t+1}$ and $p_{t+1} \ll m'_{t+1}$, consistent with (B.8). Define also (for the given non-null s_{t+1})

$$m'_t(s_{t+1}) = \frac{b'_t(s_{t+1})}{\sum_{s'_{t+1}} b'_t(s'_{t+1})} \equiv \frac{b'_t(s_{t+1})}{B'_t} > 0 \quad \text{and}$$

$$U_{t+1}(M_{t+1}, s_{t+1}) = U_{t+1}^*(M_{t+1}, s_{t+1}) / m'_t(s_{t+1}).$$

On the other hand, if s_{t+1} is \succeq_t -null, then define $m'_t(s_{t+1}) = 0$ and specify $U_{t+1}(\cdot, s_{t+1})$ to have the desired form (see the discussion above regarding null states).

Then

$$\mathcal{U}_t(F_t) = \sum_{s_{t+1}} m'_t(s_{t+1}) U_{t+1}(F_t(s_{t+1}), s_{t+1}),$$

and for every \succeq_t -non-null s_{t+1} ,

$$\begin{aligned} & U_{t+1}^*(M_{t+1}, s_{t+1}) / b'_t = U_{t+1}(M_{t+1}, s_{t+1}) / B'_t = \\ & \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ \begin{array}{l} \sum_{s_{t+2}} [p_{t+1}(s_{t+2}) + \alpha_{t+1} q_{t+1}(s_{t+2})] U_{t+2}(F_{t+1}(s_{t+2}), s_{t+2}) \\ + \frac{(1+a'_t)a''_t}{b'_t} B^*(c_{t+1}) \end{array} \right\} \\ & - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \left\{ \alpha_{t+1} \sum_{s_{t+2}} q_{t+1}(s_{t+2}) U_{t+2}(F'_{t+1}(s_{t+2}), s_{t+2}) + \frac{a'_t a''_t}{b'_t} B^*(c'_{t+1}) \right\}. \end{aligned} \tag{B.9}$$

Because B'_t is time t -measurable, one could have begun with $B'_t U_{t+1}^*(\cdot, s_{t+1})$ in place of $U_{t+1}^*(\cdot, s_{t+1})$. As a result, there is no loss of generality in assuming $B'_t = 1$ in (B.9).

We can simplify further by considering the ranking of risky consumption streams $\vec{c} = (\vec{c}_t)$, elements of \mathcal{C}^r . For these streams, (B.9) reduces (using a convenient abuse of notation and continuing to focus on \succeq_t -non-null s_{t+1}) to

$$U_{t+1}(c_{t+1}, \dots, c_T, s_{t+1}) = \frac{a''_t}{b'_t} B^*(c_{t+1}) + U_{t+2}(c_{t+2}, \dots, c_T, s_{t+2}).$$

By the inductive hypothesis,

$$U_{t+2}(c_{t+2}, \dots, c_T, s_{t+2}) = u(c_{t+2}) + \delta U_{t+3}(c_{t+3}, \dots, c_T, s_{t+3}).$$

Deduce that

$$U_{t+1}(c_{t+1}, \dots, c_T, s_{t+1}) = \frac{a''_t}{b'_t} B^*(c_{t+1}) + u(c_{t+2}) + \delta U_{t+3}(c_{t+2}, \dots, c_T, s_{t+2}),$$

and from Risk Preference conclude that wlog

$$\frac{a_t''}{b_t'} = \delta^{-1} \text{ and } B^*(\cdot) = u(\cdot).$$

Substitution into (B.9) yields

$$\begin{aligned} & \delta U_{t+1}(M_{t+1}, s_{t+1}) = \\ & \max_{(c_{t+1}, F_{t+1}) \in M_{t+1}} \left\{ \delta (1 + \alpha_{t+1}) \Sigma_{s_{t+2}} \frac{p_{t+1}(s_{t+2}) + \alpha_{t+1} q_{t+1}(s_{t+2})}{1 + \alpha_{t+1}} U_{t+2}(F_{t+1}, s_{t+2}) \right. \\ & \quad \left. + (1 + a_t') u(c_{t+1}) \right\} \\ & - \max_{(c'_{t+1}, F'_{t+1}) \in M_{t+1}} \left\{ \delta \alpha_{t+1} \Sigma_{s_{t+2}} q_{t+1}(s_{t+2}) U_{t+2}(F'_{t+1}, s_{t+2}) + a_t' u(c'_{t+1}) \right\}. \end{aligned} \tag{B.10}$$

Wlog replace δU_{t+1} by U_{t+1} on the left side of the equation.

To complete the inductive step, it remains to prove that $\alpha_{t+1}(s_{t+1}) = a_t'(s_{t+1})$. This is done by invoking Invariant Discounting: fix s_{t+1} (\succeq_t -non-null) and suppose that

$$\alpha_{t+1}(s_{t+1}) \neq a_t'(s_{t+1}). \tag{B.11}$$

Then at least one is nonzero and hence there exist contingent menus satisfying

$$(F_{-s_{t+1}}, (c, F')) \succ_t (F_{-s_{t+1}}, \{(c, F'), (d, \vec{d})\}) \succ_t (F_{-s_{t+1}}, (d, \vec{d})).$$

The corresponding strict rankings are valid also for all (e, \vec{e}) in an open neighborhood of (d, \vec{d}) . This implies the hypothesis (3.8) of Invariant Discounting. Then, by (B.10), the equivalence (3.9) translates into the statement that, for all such (e, \vec{e}) ,

$$\begin{aligned} \delta \alpha_{t+1} U_{t+2}^e + a_t' u(e) & \leq \delta \alpha_{t+1} U_{t+2}^d + a_t' u(d) \iff \\ \delta U_{t+2}^e + u(e) & \leq \delta U_{t+2}^d + u(d) \end{aligned}$$

where $\vec{d} = (d_{t+2}, \dots)$, $U_{t+2}^d = \sum_{\tau=t+2}^T \delta^{\tau-t-2} u(d_\tau)$ and similarly for U_{t+2}^e . This contradicts (B.11).

This completes the inductive step. It remains to establish the appropriate representation for preference at $T - 2$. This amounts essentially to the 3-period result in [5]. One difference is that only terminal consumption is permitted there, but the argument there is readily adapted to accommodate intermediate consumption. Another point where extra care must be taken in the present setting is in the

specification of terminal utility $U_T^*(\cdot, s_T)$. The latter is evaluated on consumption streams, which involve only risk. Thus by Risk Preference, $U_T^*(\cdot, s_T)$ is ordinal, and hence also cardinally, equivalent to $\hat{u}(\cdot) = \Sigma \delta^r u(\cdot)$ on \mathcal{C}^r . Ignoring irrelevant additive terms, we can write

$$U_T^*(\cdot, s_T) = \mu_T(s_T) \hat{u}(\cdot) \text{ on } \mathcal{C}^r.$$

This gives the proper decomposition of terminal payoffs into ‘taste’ and a state-dependent term that can be transformed into a probability. (In [5], terminal payoffs were given by $\hat{u}(\cdot)$ alone.)

Step 5(ii): Repeat the preceding on $D_{T+1} \supset D_T$. Verify that the construction of p_0 and $(p_t, q_t, \alpha_t)_1^T$ can be carried out without changing p_0 and $(p_t, q_t, \alpha_t)_1^{T-1}$ constructed as above using D_T . Proceeding in this way, one obtains p_0 and $(p_t, q_t, \alpha_t)_1^\infty$ such that each \succeq_t has the representation (2.6), (2.7) and (B.8) on $\cup_1^\infty D_T$. But the latter is dense in \mathcal{C} (Theorem A.1). Apply Continuity to obtain the desired representation on all of \mathcal{C} . (To elaborate, the preceding yields (for each t) the function $U_{t+1}(\cdot; s_{t+1})$ defined on $\cup_{\tau \geq t+1} \mathcal{M}(C \times D_\tau)$ and representing $\succeq_{t|s_{t+1}}$ there. Moreover, $U_{t+1}(\cdot; s_{t+1})$ is bounded above by $(1 - \delta)^{-1} \max_{c \in C} u(c)$ and below by $(1 - \delta)^{-1} \min_{c \in C} u(c)$. An elementary fact is: If Y is compact metric, Y' is a dense subset, \succeq is a complete, transitive and continuous order on Y and $U : Y' \rightarrow \mathbb{R}^1$ is bounded and represents \succeq restricted to Y' , then U can be extended uniquely to a real-valued function on Y so as to represent \succeq there; moreover, $U(y) = \lim U(y'_n)$ for any sequence $\{y'_n\}$ in Y' that converges to y . Apply this fact to deduce that $U_{t+1}(\cdot; s_{t+1})$ can be extended uniquely to $\mathcal{M}(C \times \mathcal{C})$ so as to represent $\succeq_{t|s_{t+1}}$ there and so as to satisfy the recursive relation (2.7).)

Step 6: Apply Bias Persistence to show that

$$m'_{t+1} = \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}}. \quad (\text{B.12})$$

The representation derived thus far has the form:

$$\mathcal{U}_t(F) = \delta \int_{S_{t+1}} U_{t+1}(F(s_{t+1}), s_{t+1}) dm'_t(s_{t+1}), \quad F \in \mathcal{C},$$

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

$$\max_{(c, F) \in M} (1 + \alpha_{t+1}) \left\{ u(c) + \delta \int_{S_{t+2}} U_{t+2}(F(s_{t+2}), s_{t+2}) d \frac{p_{t+1} + \alpha_{t+1} q_{t+1}}{1 + \alpha_{t+1}} \right\}$$

$$- \max_{(c', F') \in M} \alpha_{t+1} \left\{ u(c') + \delta \int_{S_{t+2}} U_{t+2}(F'(s_{t+2}), s_{t+2}) dq_{t+1}(s_{t+2}) \right\}.$$

Case 1: Suppose that at $t + 1$ and history s_1^{t+1} ,

$$\alpha_{t+1} = 0 \text{ or } q_{t+1}(\cdot) = p_{t+1}(\cdot).$$

Then $U_{t+1}(\cdot, s_{t+1})$ is strategically rational and thus, for any F' and F ,

$$\begin{aligned} \{(c, F')\} \succ_{t|s_{t+1}} \{(c, G)\} &\implies \\ \{(c, F')\} \sim_{t|s_{t+1}} \{(c, F'), (c, G)\} &\implies \\ F' \succeq_{t+1} G. \end{aligned}$$

where the last implication uses Bias Persistence. In particular, if F' and G lie in $\mathcal{C}_{+1}^0 = D_1$, then the above translates into the implication

$$\begin{aligned} \int_{S_{t+2}} (\hat{u}(F'(s_{t+2})) - \hat{u}(G(s_{t+2}))) dp_{t+1} > 0 &\implies \\ \int_{S_{t+2}} (\hat{u}(F'(s_{t+2})) - \hat{u}(G(s_{t+2}))) dm'_{t+1} \geq 0, \end{aligned}$$

where $\hat{u}(G(s_{t+2}))$ is the utility, computed using $\Sigma \delta^r u(\cdot)$, of the consumption stream implied by $G(s_{t+2})$ and so on. It follows from a Theorem of the Alternative [15, p. 34] that $m'_{t+1} = p_{t+1}$, which implies (B.12).

Case 2: Suppose on the other hand that

$$\alpha_{t+1} > 0 \text{ and } q_{t+1}(\cdot) \neq p_{t+1}(\cdot).$$

Then $p_{t+1} + \alpha_{t+1}q_{t+1} \neq q_{t+1}$ and hence there exist G and H in \mathcal{C} such that

$$\begin{aligned} \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1}q_{t+1}] > \int_{S_{t+2}} U_{t+2}(H(s_{t+2}), s_{t+2}) d[p_{t+1} + \alpha_{t+1}q_{t+1}] \quad \text{and} \\ \int_{S_{t+2}} U_{t+2}(G(s_{t+2}), s_{t+2}) dq_{t+1} < \int_{S_{t+2}} U_{t+2}(H(s_{t+2}), s_{t+2}) dq_{t+1}. \end{aligned} \tag{B.13}$$

Let F' be sufficiently close to G so that these inequalities are preserved if G is replaced by F' . Take any c in C and F in \mathcal{C} . Then (3.11) is satisfied iff

$$\int_{S_{t+2}} (U_{t+2}(F'(s_{t+2}), s_{t+2}) - U_{t+2}(G(s_{t+2}), s_{t+2})) d[p_{t+1} + \alpha_{t+1}q_{t+1}] \geq 0. \quad (\text{B.14})$$

Thus Bias Persistence implies that for all G and H satisfying (B.13), and F' sufficiently close to G , then (B.14) implies

$$\int_{S_{t+2}} (U_{t+2}(F'(s_{t+2}), s_{t+2}) - U_{t+2}(G_{t+1}(s_{t+2}), s_{t+2})) dm'_{t+1} \geq 0.$$

But the integrand in the latter two inequalities varies over an open neighborhood of zero as F' , G and H vary over the conditions noted above. Thus (B.12) follows by a Theorem of the Alternative [15, p. 34] ■

C. APPENDIX: Proofs for Specific Biases

Proof of Corollary 4.1: Necessity of Prior-Bias: Given the representation, let

$$\mu_t(s_{t+2}) = \int p_{t+1}(s_{t+2} | s'_{t+1}) dm_t(s'_{t+1}).$$

Then the axiom can be translated into the statement:

$$\begin{aligned} \int [U_{t+2}(F(s_{t+2}), s_{t+2}) - U_{t+2}(G(s_{t+2}), s_{t+2})] dp_{t+1} &> 0 \quad \text{and} \\ \int [U_{t+2}(F(s_{t+2}), s_{t+2}) - U_{t+2}(G(s_{t+2}), s_{t+2})] d\mu_t(s_{t+2}) &= 0 \end{aligned}$$

imply

$$\alpha_{t+1} \int [U_{t+2}(F(s_{t+2}), s_{t+2}) - U_{t+2}(G(s_{t+2}), s_{t+2})] dq_{t+1} \geq 0.$$

This is obviously satisfied given (4.6).

Sufficiency of Prior-Bias: If $\alpha_{t+1} = 0$, then

$$\alpha_{t+1} (q_{t+1}(\cdot | s_{t+1}) - p_{t+1}(\cdot | s_{t+1})) = 0$$

and we can take any q_{t+1} , including in particular q_{t+1} as in (4.6) with any λ_{t+1} .

Suppose that $\alpha_{t+1} > 0$. Then Prior-Bias asserts the implication described above, where $\alpha_{t+1} = 1$ wlog. To argue further, eliminate the state dependence in U_{t+2} by supposing that both F and G lie in \mathcal{C}_{+1}^0 , that is, they provide perfect commitment and that all uncertainty is resolved at $t + 2$. Then, recalling the discussion surrounding (2.4), $F(s_{t+2})$ and $G(s_{t+2})$ can be identified with deterministic consumption process $c^F(s_{t+2})$ and $c^G(s_{t+2})$ respectively. It follows that

$$U_{t+2}(F(s_{t+2}), s_{t+2}) = \Sigma_{\tau=t+2} \delta^{\tau-(t+2)} u(c_{\tau}^F(s_{t+2})) \equiv \widehat{u}(c^F(s_{t+2})),$$

and similarly for G . Write

$$x^{FG}(s_{t+2}) = \widehat{u}(c^F(s_{t+2})) - \widehat{u}(c^G(s_{t+2})).$$

Then

$$\left[\int x^{FG}(s_{t+2}) dp_{t+1}(s_{t+2} | s_{t+1}) > 0 \text{ and } \int x^{FG}(s_{t+2}) d\mu_t(s_{t+2}) = 0 \right] \implies \int x^{FG}(s_{t+2}) dq_{t+1} \geq 0.$$

One can show that $x^{FG}(\cdot)$ can be made to vary sufficiently (over an open neighborhood of zero) as we range over F and G lying in \mathcal{C}_{+1}^0 . Apply a Theorem of the Alternative [15, p. 34].

The arguments for the other axioms are similar. ■

Proof of Corollary 4.2: The proof is similar to that of the preceding corollary. We point out only that for G in \mathcal{C}_{+1}^0 , $U_{t+1}(c_{t+1}, (\int G d\Psi_{t+1})(s_{t+1}), s_{t+1})$

$$\begin{aligned} &= u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2} \left(\int_{S_{t+2}} G(s'_{t+2}) d\Psi_{t+1, s_{t+2}} \right) dp_{t+1}(s_{t+2}) \\ &= u(c_{t+1}) + \delta \int_{S_{t+2}} \int_{S_{t+2}} U_{t+2}(G(s'_{t+2}), s_{t+2}) d\Psi_{t+1}(s'_{t+2}) dp_{t+1}(s_{t+2}) \\ &= u(c_{t+1}) + \delta \int_{S_{t+2}} U_{t+2}(G(s'_{t+2}), s_{t+2}) d\Psi_{t+1}(s'_{t+2}), \end{aligned}$$

because $U_{t+2}(G(s'_{t+2}), s_{t+2})$ does not depend on s_{t+2} . ■

D. APPENDIX: Learning in the Long Run

Proof of Theorem 5.1: (a) First we show that $\log \mu_t(\theta^*)$ is a submartingale under P^* . Because

$$\log \mu_{t+1}(\theta^*) - \log \mu_t(\theta^*) = \log \left((1 - \gamma_{t+1}) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right), \quad (\text{D.1})$$

it suffices to show that

$$E^* \left[\log \left((1 - \gamma_{t+1}) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \mid \mathcal{S}_t \right] \geq 0, \quad (\text{D.2})$$

where E^* denotes expectation with respect to P^* . By assumption, γ_{t+1} is constant given \mathcal{S}_t . Thus the expectation equals

$$\begin{aligned} & \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) \log \left((1 - \gamma_{t+1}) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \geq \\ & \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) (1 - \gamma_{t+1}) \log \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) = \\ & (1 - \gamma_{t+1}) \sum_{s_{t+1}} \ell(s_{t+1} \mid \theta^*) \log \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \right) \geq 0 \end{aligned}$$

as claimed, where both inequalities are due to concavity of $\log(\cdot)$. (The second is the well-known entropy inequality.)

Clearly $\log \mu_t(\theta^*)$ is bounded above by zero. Therefore, by the martingale convergence theorem, it converges $P^* - a.s.$ From (D.1),

$$\log \mu_{t+1}(\theta^*) - \log \mu_t(\theta^*) = \log \left((1 - \gamma_{t+1}) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \longrightarrow 0$$

and hence $\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \longrightarrow 1$ $P^* - a.s.$ ■

(b) $E^* \left[\left((1 - \gamma_{t+1}) \frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} + \gamma_{t+1} \right) \mid \mathcal{S}_t \right] = (1 - \gamma_{t+1}) E^* \left[\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} \mid \mathcal{S}_t \right] + \gamma_{t+1} \geq (1 - \gamma_{t+1}) + \gamma_{t+1} = 1$. (The last inequality is implied by the fact that

$$\min_X \left\{ E^* \left[\frac{1}{X(s_{t+1})} \mid \mathcal{S}_t \right] : E^* [X(s_{t+1}) \mid \mathcal{S}_t] = 1 \right\} = 1.$$

The minimization is over random variable X 's, $X : S_{t+1} \longrightarrow \mathbb{R}_{++}^1$, and it is achieved at $X(\cdot) = 1$ because $\frac{1}{x}$ is a convex function on $(0, \infty)$.) Deduce that

$E^* \left[\frac{\mu_{t+1}(\theta^*)}{\mu_t(\theta^*)} \mid \mathcal{S}_t \right] \geq 1$ and hence that $\mu_t(\theta^*)$ is a submartingale. By the martingale convergence theorem,

$$\mu_\infty(\theta^*) \equiv \lim \mu_t(\theta^*) \quad \text{exists } P^* - a.s.$$

Claim: $\mu_\infty(\theta^*) > 0$ on a set with positive P^* -probability: By the bounded convergence theorem,

$$E^* \mu_t(\theta^*) \longrightarrow E^* \mu_\infty(\theta^*);$$

and $E^* \mu_t(\theta^*) \nearrow$ because $\mu_t(\theta^*)$ is a submartingale. Thus $\mu_0(\theta^*) > 0$ implies that $E^* \mu_\infty(\theta^*) > 0$, which proves the claim.

It suffices now to show that if $\mu_\infty(\theta^*) > 0$ along a sample path s_1^∞ , then forecasts are eventually correct along s_1^∞ . But along such a path, $\frac{\mu_{t+1}(\theta^*)}{\mu_t(\theta^*)} \longrightarrow 1$ and hence

$$(1 - \gamma_{t+1}) \left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} - 1 \right) \longrightarrow 0.$$

By assumption, $(1 - \gamma_{t+1})$ is bounded away from zero. Therefore,

$$\left(\frac{\ell(s_{t+1}|\theta^*)}{m_t(s_{t+1})} - 1 \right) \longrightarrow 0. \quad \blacksquare$$

Part (c) calls for two examples.

Example 1: Convergence to wrong forecasts may occur with P^* -positive probability when $\gamma_{t+1} < 0$, even where γ_{t+1} is \mathcal{S}_t -measurable (overreaction); in fact, we take $(\alpha_{t+1}, \lambda_{t+1}) = (\alpha, \lambda)$ and hence also $\gamma_{t+1} = \gamma$ to be constant over time and states.

Think of repeatedly tossing an unbiased coin that is viewed at time 0 as being either unbiased or having probability of Heads equal to b , $0 < b < \frac{1}{2}$. Thus take $S = \{H, T\}$ and $\ell(H | \theta) = \theta$ for $\theta \in \Theta = \{b, \frac{1}{2}\}$. Assume also that

$$1 < -\gamma < \frac{b}{\frac{1}{2} - b}. \quad (\text{D.3})$$

The inequality $\gamma < -1$ indicates a sufficient degree of overreaction.

To explain the reason for the other inequality, note that the model requires that (ν_t) solving (5.6) be a probability measure (hence non-negative valued). This is trivially true if $\lambda_{t+1} \geq 0$ but otherwise requires added restrictions: $\nu_{t+1} \geq 0$ if

$$\frac{\ell(s_{t+1} | \theta)}{m_t(s_{t+1})} = \frac{dBU(\mu_t; s_{t+1})(\theta)}{d\mu_t} \geq -\frac{\lambda_{t+1}}{1 + \lambda_{t+1}}.$$

In the present example $\min_{s, \theta} \frac{\ell(s|\theta)}{m_t(s)} \geq 2b$, and thus it suffices to have

$$-\frac{\lambda}{1 + \lambda} \leq 2b. \quad (\text{D.4})$$

Because any non-negative value for α is admissible, $\gamma = \frac{\alpha\lambda}{1+\alpha}$ is consistent with (D.4) if and only if $-\gamma < b/(\frac{1}{2} - b)$.

We show that if (D.3), then

$$m_t(\cdot) \longrightarrow \ell(\cdot | b) \quad \text{as } t \longrightarrow \infty,$$

with probability under P^* at least $\frac{1}{2}$.

Abbreviate $\mu_t(\frac{1}{2})$ by μ_t^* .

Claim 1: $\mu_\infty^* \equiv \lim \mu_t^*$ exists $P^* - a.s.$ and if $\mu_\infty^* > 0$ for some sample realization s_1^∞ , then $m_t(H) \longrightarrow \frac{1}{2}$ and $\mu_t^* \longrightarrow 1$ along s_1^∞ . (The proof is analogous to that of part (b).) Deduce that

$$\mu_\infty^* \in \{0, 1\} \quad P^* - a.s.$$

Claim 2: $f(z) \equiv \left[(1 - \gamma) \frac{1}{z} + \gamma \right] \left[(1 - \gamma) \frac{1 - \frac{1}{2}}{(1 - z)} + \gamma \right] \leq 1$, for all $z \in [b, \frac{1}{2}]$. Argue that $f(z) \leq 1 \iff g(z) \equiv [(1 - \gamma) + 2\gamma z] [(1 - \gamma) + 2\gamma(1 - z)] - 4z(1 - z) \leq 0$. Compute that $g(\frac{1}{2}) = 0$, $g'(\frac{1}{2}) = 0$ and g is concave because $\gamma < -1$. Thus $g(z) \leq g(0) = 0$.

Claim 3: $E^* \left[\log \left((1 - \gamma) \frac{\ell(s_{t+1} | \frac{1}{2})}{m_t(s_{t+1})} + \gamma \right) \mid \mathcal{S}_t \right] = \frac{1}{2} \log \left((1 - \gamma) \frac{1}{b + (\frac{1}{2} - b)\mu_t^*} + \gamma \right) + \frac{1}{2} \log \left((1 - \gamma) \frac{1 - \frac{1}{2}}{(1 - b - (\frac{1}{2} - b)\mu_t^*)} + \gamma \right) = \frac{1}{2} \log (f(b + (\frac{1}{2} - b)\mu_t(\frac{1}{2}))) \leq 0$, by Claim 2.

By Claim 1, it suffices to prove that $\mu_\infty^* = 1$ $P^* - a.s.$ is impossible. Compute that

$$\mu_t^* = \mu_0^* \left[\prod_{k=0}^{t-1} \left((1 - \gamma) \frac{\ell(s_{k+1} | \frac{1}{2})}{m_k(s_{k+1})} + \gamma \right) \right],$$

$$\begin{aligned} \log \mu_t^* &= \log \mu_0^* + \sum_{k=0}^{t-1} \log \left((1 - \gamma) \frac{\ell(s_{k+1} | \frac{1}{2})}{m_k(s_{k+1})} + \gamma \right) \\ &= \log \mu_0^* + \sum_{k=0}^{t-1} (\log z_{k+1} - E[\log z_{k+1} | \mathcal{S}_k]) + \sum_{k=0}^{t-1} E[\log z_{k+1} | \mathcal{S}_k], \end{aligned}$$

where $z_{k+1} = (1 - \gamma) \frac{\ell(s_{k+1} | \frac{1}{2})}{m_k(s_{k+1})} + \gamma$. Therefore, $\log \mu_t^* \geq \frac{1}{2} \log \mu_0^*$ iff

$$\sum_{k=0}^{t-1} (\log z_{k+1} - E[\log z_{k+1} | \mathcal{S}_k]) \geq -\frac{1}{2} \log \mu_0^* - \sum_{k=0}^{t-1} E[\log z_{k+1} | \mathcal{S}_k] \equiv a_k.$$

By Claim 3, $a_k > 0$. The random variable $\log z_{k+1} - E[\log z_{k+1} | \mathcal{S}_k]$ takes on two possible values, corresponding to $s_{k+1} = H$ or T , and under the truth they are equally likely and average to zero. Thus

$$P^*(\log z_{k+1} - E[\log z_{k+1} | \mathcal{S}_k] \geq a_k) \leq \frac{1}{2}.$$

Deduce that

$$P^*(\log \mu_t^* \geq \frac{1}{2} \log \mu_0^*) \leq \frac{1}{2}$$

and hence that

$$P^*(\log \mu_t^* \rightarrow 0) \leq \frac{1}{2}. \quad \blacksquare$$

Example 2: Convergence to wrong forecasts may occur with P^* -positive probability when $\gamma_{t+1} > 0$ (Positive Prior-Bias), if γ_{t+1} is only \mathcal{S}_{t+1} -measurable.

The coin is as before - it is unbiased, but the agent does not know that and is modeled via $S = \{H, T\}$ and $\ell(H | \theta) = \theta$ for $\theta \in \Theta = \{b, \frac{1}{2}\}$. Assume further that α_{t+1} and λ_{t+1} are such that

$$\gamma_{t+1} \equiv \frac{\alpha_{t+1} \lambda_{t+1}}{1 + \alpha_{t+1}} = \begin{cases} w & \text{if } s_{t+1} = H \\ 0 & \text{if } s_{t+1} = T, \end{cases}$$

where $0 < w < 1$. Thus, from (5.9), the agent updates by Bayes' Rule when observing T but attaches only the weight $(1 - w)$ to last period's prior when observing H . Assume that

$$w > 1 - 2b.$$

Then

$$m_t(\cdot) \rightarrow \ell(\cdot | b) \quad \text{as } t \rightarrow \infty,$$

with probability under P^* at least $\frac{1}{2}$.

The proof is similar to that of Example 1. The key is to observe that $E^* \left[\log \left((1 - \gamma) \frac{\ell(s_{t+1}|\frac{1}{2})}{m_t(s_{t+1})} + \gamma \right) \mid \mathcal{S}_t \right] \leq 0$ under the stated assumptions.

The proof of Theorem 5.2 requires the following lemmas:

Lemma D.1. (Freedman (1975)) *Let $\{z_t\}$ be a sequence of uniformly bounded \mathcal{S}_t -measurable random variables such that for every $t \geq 1$, $E^*(z_{t+1}|\mathcal{S}_t) = 0$. Let $V_t^* \equiv \text{VAR}(z_{t+1}|\mathcal{S}_t)$ where VAR is the variance operator associated with P^* . Then,*

$$\sum_{t=1}^n z_t \text{ converges to a finite limit as } n \rightarrow \infty, P^*\text{-a.s. on } \left\{ \sum_{t=1}^{\infty} V_t^* < \infty \right\}$$

and

$$\sup_n \sum_{t=1}^n z_t = \infty \text{ and } \inf_n \sum_{t=1}^n z_t = -\infty, P^*\text{-a.s. on } \left\{ \sum_{t=1}^{\infty} V_t^* = \infty \right\}.$$

Definition D.2. *A sequence of $\{x_t\}$ of \mathcal{S}_t -measurable random variables is eventually a submartingale if, $P^* - a.s.$, $E^*(x_{t+1}|\mathcal{S}_t) - x_t$ is strictly negative at most finitely many times.*

Lemma D.3. *Let $\{x_t\}$ be uniformly bounded and eventually a submartingale. Then, $P^* - a.s.$, x_t converges to a finite limit as t goes to infinity.*

Proof. Write

$$x_t = \sum_{j=1}^t (r_j - E^*(r_j|\mathcal{S}_{j-1})) + \sum_{j=1}^t E^*(r_j|\mathcal{S}_{j-1}) + x_0, \text{ where } r_j \equiv x_j - x_{j-1}.$$

By assumption, $P^* - a.s.$, $E^*(r_j|\mathcal{S}_{j-1})$ is strictly negative at most finitely many times. Hence, $P^* - a.s.$,

$$\inf_t \sum_{j=1}^t E^*(r_j|\mathcal{S}_{j-1}) > -\infty.$$

Given that x_t is uniformly bounded, $P^* - a.s.$,

$$\sup_t \sum_{j=1}^t z_j < \infty, \text{ where } z_j \equiv r_j - E^*(r_j | \mathcal{S}_{j-1}).$$

It follows from Freedman's result that $P^* - a.s.$,

$$\sum_{j=1}^t z_j \text{ converges to a finite limit as } t \rightarrow \infty.$$

It now follows from x_t uniformly bounded that $\sup_t \sum_{j=1}^t E^*(r_j | \mathcal{S}_{j-1}) < \infty$. Because $E^*(r_j | \mathcal{S}_{j-1})$ is strictly negative at most finitely many times,

$$\sum_{j=1}^t E^*(r_j | \mathcal{S}_{j-1}) \text{ converges to a finite limit as } t \rightarrow \infty.$$

Therefore, $P^* - a.s.$, x_t converges to a finite limit as t goes to infinity. ■

Proof of Theorem 5.2:

Claim 1: Define $f(\theta, m) = \sum_k \theta_k^* \frac{\theta_k}{m_k}$ on the interior of the $2K$ -simplex. There exists $\delta' \in \mathbb{R}_{++}^K$ such that

$$|\theta_k - \theta_k^*| < \delta'_k \text{ for all } k \implies f(\theta, m) - 1 \geq -\underline{\gamma} K^{-1} \sum_k |m_k - \theta_k|.$$

Proof: $f(\theta, \theta) = 1$, $f(\theta, \cdot)$ is convex and hence

$$\begin{aligned} f(\theta, m) - 1 &\geq \sum_{k \neq K} \left(\frac{\partial f(\theta, m)}{\partial m_k} - \frac{\partial f(\theta, m)}{\partial m_K} \right) \Big|_{m=\theta} (m_k - \theta_k) \\ &= \sum_{k \neq K} \left(-\frac{\theta_k^*}{\theta_k} + \frac{\theta_K^*}{\theta_K} \right) (m_k - \theta_k). \end{aligned}$$

But the latter sum vanishes at $\theta = \theta^*$. Thus argue by continuity.

Given any $\delta \in \mathbb{R}_{++}^K$, $\delta \ll \delta'$, define $\Theta^* = (\theta^* - \delta, \theta^* + \delta) \equiv \prod_{k=1}^K (\theta_k^* - \delta_k, \theta_k^* + \delta_k)$ and $\mu_t^* = \sum_{\theta \in \Theta^*} \mu_t(\theta)$.

Claim 2: Define $m_t^*(s^k) = \sum_{\theta \in \Theta^*} \theta_k \mu_t(\theta) / \mu_t^*(\theta)$. Then

$$|m_t(s^k) - m_t^*(s^k)| \leq 1 - \mu_t^*.$$

Proof: $m_t(s^k) - m_t^*(s^k) = \frac{\sum_{\theta \in \Theta^*} \theta_k \mu_t(\theta)}{\mu_t^*} (\mu_t^* - 1) + \sum_{\theta \notin \Theta^*} \theta_k \mu_t(\theta)$. Therefore,
 $(\mu_t^* - 1) \leq m_t^*(s^k) (\mu_t^* - 1) = \frac{\sum_{\theta \in \Theta^*} \theta_k \mu_t(\theta)}{\mu_t^*} (\mu_t^* - 1) \leq m_t(s^k) - m_t^*(s^k) \leq \sum_{\theta \notin \Theta^*} \theta_k \mu_t(\theta) \leq 1 - \mu_t^*$.

Claim 3: For any $\delta \ll \delta'$ as above,

$$\sum_k \theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} - 1 \geq -\underline{\gamma} (1 - \mu_t^*).$$

Proof: Because $|m_t^*(s^k) - \theta_k^*| < \delta_k < \delta'_k$, we have that

$$\sum_k \theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} - 1 \geq -\underline{\gamma} K^{-1} \sum_k |m_t(s^k) - m_t^*(s^k)|.$$

Now Claim 3 follows from Claim 2.

Compute that

$$E^* [\mu_{t+1}(\theta) | \mathcal{S}_t] = (1 - \gamma_{t+1}) \left[\sum_k \theta_k^* \frac{\theta_k}{m_t(s^k)} \right] \mu_t(\theta) + \gamma_{t+1} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t], \quad (\text{D.5})$$

where use has been made of the assumption that γ_{t+1} is \mathcal{S}_t -measurable. Therefore,

$$\begin{aligned} E^* [\mu_{t+1}^*(\theta) | \mathcal{S}_t] - \mu_t^* &= (1 - \gamma_{t+1}) \sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) \mu_t^* + \gamma_{t+1} \sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] - \mu_t^* \\ &= (1 - \gamma_{t+1}) \left[\sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) - 1 \right] \mu_t^* + \gamma_{t+1} \sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] - \gamma_{t+1} \mu_t^*. \end{aligned}$$

By the LLN, $P^* - a.s.$ for large enough t the frequency of s^k will eventually be θ_k^* and

$$\sum_{\theta \in \Theta^*} E^* [\psi_{t+1}(\theta) | \mathcal{S}_t] = 1.$$

Eventually along any such path,

$$E^* [\mu_{t+1}^*(\theta) | \mathcal{S}_t] - \mu_t^* = (1 - \gamma_{t+1}) \left[\sum_k \left(\theta_k^* \frac{m_t^*(s^k)}{m_t(s^k)} \right) - 1 \right] \mu_t^* + \gamma_{t+1} (1 - \mu_t^*)$$

$$\geq [-\underline{\gamma}(1 - \gamma_{t+1})\mu_t^* + \gamma_{t+1}](1 - \mu_t^*) \geq 0,$$

where the last two inequalities follow from Claim 3 and the hypothesis $\underline{\gamma} \leq \gamma_{t+1}$.

Hence (μ_t^*) is eventually a P^* -submartingale. By Lemma D.3, $\mu_\infty^* \equiv \lim \mu_t^*$ exists $P^* - a.s.$ Consequently, $E^*[\mu_{t+1}^*(\theta) | \mathcal{S}_t] - \mu_t^* \rightarrow 0$ $P^* - a.s.$ and from the last displayed equation, $[-\underline{\gamma}(1 - \gamma_{t+1})\mu_t^* + \gamma_{t+1}](1 - \mu_t^*) \rightarrow 0$ $P^* - a.s.$ It follows that $\mu_\infty^* = 1$. Finally, $m_t(\cdot) = \int \ell(\cdot | \theta) d\mu_t$ eventually remains in $\Theta^* = (\theta^* - \delta, \theta^* + \delta)$.

Above δ is arbitrary. Apply the preceding to $\delta = \frac{1}{n}$ to derive a set Ω_n such that $P^*(\Omega_n) = 1$ and such that for all paths in Ω_n , m_t eventually remains in $(\theta^* - \frac{1}{n}, \theta^* + \frac{1}{n})$. Let $\Omega \equiv \bigcap_{n=1}^{\infty} \Omega_n$. Then, $P^*(\Omega) = 1$ and for all paths in Ω , m_t converges to θ^* . ■

References

- [1] C.D. Aliprantis and K. Border, *Infinite Dimensional Analysis*, Springer Verlag, 1994.
- [2] A. Brandenburger and E. Dekel, Hierarchies of beliefs and common knowledge, *J. Econ. Theory* 59 (1993), 189-198.
- [3] C. Camerer, Individual decision-making, in *Handbook of Experimental Economics*, J. Kagel and A. Roth eds., Princeton U. Press, 1995.
- [4] L.G. Epstein, Stationary cardinal utility and optimal growth under uncertainty, *J. Econ. Theory* 31 (1983), 133-152.
- [5] L.G. Epstein, An axiomatic model of non-Bayesian updating, revised March 2004.
- [6] L.G. Epstein and M. Peters, A revelation principle for competing mechanisms, *J. Econ. Theory* 88 (1999), 119-160.
- [7] L.G. Epstein and T. Wang, 'Beliefs about beliefs' without probabilities, *Econometrica* 64 (1995), 1343-1373.
- [8] D. Freedman, "On tail probabilities for martingales," *Annals of Probability* 3 (1975), 100-118.

- [9] F. Gul and W. Pesendorfer, Temptation and self-control, *Econometrica* 69 (2001), 1403-35.
- [10] F. Gul and W. Pesendorfer, Self-control and the theory of consumption, forthcoming in *Econometrica*.
- [11] L. Kogan, S. Ross, J. Wang and M. Westerfield, The price impact and survival of irrational traders," MIT 2003.
- [12] D. Kreps, *Notes on the Theory of Choice*, Westview, 1988.
- [13] D. Kreps, Static choice in the presence of unforeseen contingencies, in *Essays in Honour of F. Hahn*, P. Dasgupta et al eds., MIT Press, 1992.
- [14] E. Lehrer and R. Smorodinsky, Merging and learning, *Statistics, Probability and Game Theory, IMS Lecture Notes* 30 (1996), 147-168.
- [15] O.L. Mangasarian, *Nonlinear Programming*, McGraw-Hill, New York: 1969.
- [16] J.F. Mertens and S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory*, 14 (1985), 1-29.
- [17] S. Mullainathan, Thinking through categories, MIT, 2000.
- [18] M. Rabin, Psychology and economics, *J. Ec. Lit.* 36 (1998), 11-46.
- [19] M. Rabin, Inference by believers in the law of small numbers, *Quart. J. Econ.* 117 (2002), 775-816.
- [20] A. Sandroni, Efficient markets and Bayes' Rule, NW, 2003.
- [21] R.H. Strotz, Myopia and inconsistency in dynamic utility maximization, *Rev. Ec. Studies* 23 (1955), 165-180.
- [22] A. Tversky and D. Kahneman, Belief in the law of small numbers, *Psychology Bulletin* 76 (1971), 105-110.
- [23] A. Tversky and D. Kahneman, Judgement under uncertainty: heuristics and biases, *Science* 185 (1974), 1124-31.