AN AXIOMATIC MODEL OF NON-BAYESIAN UPDATING*

Larry G. Epstein

March 2004

Abstract

This paper models an agent in a three-period setting who does not update according to Bayes' Rule, and who is self-aware and anticipates her updating behavior when formulating plans. The agent is rational in the sense that her dynamic behavior is derived from a stable preference order on a domain of state-contingent menus of acts. A representation theorem generalizes (the dynamic version of) Anscombe-Aumann's theorem so that *both* the prior *and* the way in which it is updated are subjective. The model can generate updating biases analogous to those observed by psychologists.

*Department of Economics, University of Rochester, Rochester, NY 14627, lepn@troi.cc.rochester.edu. I have benefitted from comments by David Easley, Massimo Marinacci, Jawaad Noor, Francesco Squintani, seminar participants at Cornell, UC Irvine, UCSD and Stanford, and especially from Mark Machina and Uzi Segal. I am grateful also to Igor Kopylov for expert research assistance.

1. INTRODUCTION

1.1. Motivation and Outline

This paper models an agent in a three-period setting who does not update according to Bayes' Rule, and who is self-aware and anticipates her updating behavior when formulating plans. The major contribution is a representation theorem for suitably defined preference that provides (in a sense qualified in the concluding section) axiomatic foundations for non-Bayesian updating. One perspective on the theorem is obtained through its relation to a dynamic version of the Anscombe-Aumann theorem which provides foundations for reliance on a probability measure representing subjective prior beliefs and for subsequent Bayesian updating of the prior.[1] Thus, while beliefs are subjective and can vary with the agent, updating behavior cannot - everyone must update by Bayes' Rule. This Anscombe-Aumann theorem is generalized here so as to render it more fully subjective - *both* the prior *and* the way in which it is updated are subjective.

Non-Bayesian updating leads to changing beliefs and hence to changing preferences over alternatives (Anscombe-Aumann acts). This in turn leads to the temptation to deviate from previously formulated plans. Thus we are led to adapt the Gul and Pesendorfer [10, 11] model of temptation and self-control. While these authors (henceforth GP) strive to explain 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 0, the agent has a prior view of the relationship between the next observation s_1 and the future uncertainty s_2. But after observing a particular realization s_1, she changes her view on the noted relationship. For example, she may respond exuberantly to a good signal after it is realized and decide that it is an even better signal about future states than she had thought ex ante. Or the realization of a bad signal may lead her to panic, that is, to interpret the signal as an even worse omen for the future than she had thought ex ante. In either case, it is as though she retroactively changes her prior and then applies Bayes' Rule to the

[1] A rough statement is as follows: if conditional preference at every decision node conforms to subjective expected utility theory and is independent of unrealized parts of the tree, then preferences are dynamically consistent if and only if they have a common vNM index and conditional beliefs at each node are derived by Bayesian updating of the initial prior. For a recent formalization in a Savage-style setting see [9]. For an Anscombe-Aumann setting, which is more relevant to this paper, the assertion is a special case of the main result in [7] regarding the updating of sets of priors.

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; for example, 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 make choices 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.[2] Menus play a central role in this model as well. We model preference over *contingent menus* and show that the associated behavior reveals underlying (non-Bayesian) updating. As in GP, by assuming that preference is defined over (contingent) menus, we are able to model the agent's dynamic behavior via maximization of a stable (complete and transitive) preference relation. This is possible because our agent is sophisticated - she is forward-looking and anticipates her exuberance or, more generally, her psyche as it affects her reactions to signals ex post. Are individuals typically self-aware to this degree? We are not familiar with definitive evidence on this question and in its absence, we are inclined to feel that full self-awareness is a plausible working hypothesis.[3] Even where the opposite extreme of complete naivete seems descriptively more accurate, our model may help to clarify which economic consequences are due to non-Bayesian updating per se and which are due to naivete. Finally, it is the agent's sophistication that permits updating behavior to be inferred from her (in principle observable) ranking of contingent menus. This permits us to model "time-varying beliefs" while staying within the choice-theoretic tradition of Savage. We think it worthwhile to explore such modest departures from standard models before discarding the entire framework.

Several systematic deviations from Bayesian updating have been discussed in the psychology literature and some of these have been incorporated into modeling exercises in behavioral finance.[4] Our model cannot address these findings directly

[2] Kreps [12, 14] was the first to point out the advantage of modeling preference over menus. See Dekel, Lipman and Rustichini [6], Nehring [17], and Barbera and Grodal [2] for more recent refinements and variations.

[3] Comparable sophistication is assumed by GP and also in the literature on non-geometric discounting where the agent is often modeled as gaming against herself (see Laibson [15], for example).

[4] See the surveys [5] and [18] for references to the psychology literature. Two recent applications in finance that contain extensive bibliographies to the behavioral literature are [3, 4].

because the experimental literature deals with settings where prior probabilities are given objectively, while our model deals with the case where probabilities are subjective (which case we would argue is more relevant for economic modeling). Nevertheless, we show in Section 3.3 that our model can accommodate updating biases analogous to several discussed by psychologists and in the behavioral economics literature. This serves to demonstrate the richness of the model. Because it is also axiomatic, we suggest that it may provide a useful encompassing framework for addressing updating and related behavior.

The difference between objective and subjective probabilities is important for how one thinks about non-Bayesian updating. When probabilities are objective, deviations from Bayes' Rule are typically viewed as mistakes, the results of bounded rationality in light of the complexity and nonintuitive nature of Bayes' Rule (see Tversky and Kahneman [21, p. 1130], for example). We agree with this view when probabilities are objective. However, updating behavior can be understood differently when probabilities are subjective. As described above, the agent in our model is sophisticated and she uses Bayes' Rule, but she applies it to a retroactively changing prior that is triggered by the realization of a signal. Changing priors retroactively is not a 'mistake' or a sign of irrationality. After all, there are no 'correct' beliefs here, only an initial prior formulated ex ante and the agent is presumably entitled to change her view of the world given the new perspective afforded by the realization of a particular signal. This way of thinking of non-Bayesian updating in terms of changing priors is reminiscent of the literature, stemming from Strotz [20], concerning non-geometric discounting and changing tastes. In the latter class of models, tastes are assumed to change with the passage of time. Here, on the other hand, it is beliefs rather than tastes that change, and the change is triggered by the realization of a signal rather than by the passage of time per se.

1.2. Updating, Temptation and Self-Control

This section provides a brief outline of the GP model and the way in which we adapt it.

Let $\Delta(X)$ denote the set of lotteries with payoffs in X and let $\succeq$ be a preference relation on menus of lotteries (suitably closed subsets of $\Delta(X)$). The interpretation is that at an unmodeled ex post stage, a lottery is selected from the menu chosen ex ante according to $\succeq$. GP axiomatize a representation for $\succeq$ of the form

$$\mathcal{U}(A) = \max_{x \in A} \left\{ U(x) + V(x) - \max_{y \in A} V(y) \right\}, \tag{1.1}$$

for any menu A, where U and V are vNM utility functions over lotteries. For singleton menus, $\mathcal{U}(\{x\}) = U(x)$ and thus U describes preference under commitment. The function V describes the agent's urges at the second stage. In the absence of commitment, there is a temptation to maximize V and hence to deviate from the ex post choices that would be prescribed by U. Temptation can be resisted, but at the cost of self-control given by $\max_{y \in A} V(y) - V(x)$. A balance between commitment preference and the cost of self-control is achieved by choosing a lottery ex post that maximizes the compromise utility function $U + V$.

Temptation and self-control costs are illustrated behaviorally by the ranking

$$\{x\} \succ \{x, y\} \succ \{y\}. \tag{1.2}$$

The strict preference for $\{x\}$ over $\{x, y\}$ indicates that even though x is strictly preferred to y under commitment, the presence of y in the menu is tempting. The ranking $\{x, y\} \succ \{y\}$ reveals that self-control is exercised to resist the temptation and to choose x out of $\{x, y\}$. The above intuition is captured more generally in GP's central axiom of Set-Betweenness:[5] For all menus A and B,

$$A \succeq B \implies A \succeq A \cup B \succeq B.$$

In the special case where $A \succeq B \implies A \sim A \cup B$, a menu today can be valued according to the best lottery in the menu as in the standard approach. See Kreps [13, Ch. 13] who coins the label *strategic rationality*.

The model to follow combines key elements of the GP model with the Anscombe-Aumann model of subjective probability. At a formal level, we adapt the GP model by introducing state spaces and then considering preferences over (suitably contingent) menus of Anscombe-Aumann acts rather than over menus of lotteries. To elaborate, there are 3 periods - an ex ante stage 0, an interim period 1 when a signal $s_1 \in S_1$ is realized, and period 2 when remaining uncertainty is resolved through realization of some $s_2 \in S_2$. At time 0, the agent chooses some F, an s_1-contingent menu of acts over S_2. She does this cognizant of the fact that at

[5] The axiom restricts the nature of temptation so that a set of alternatives is as tempting as its most tempting member. See GP (pp. 1408-9) for reasons why Set-Betweenness might be violated.

time 1, after a particular s_1 is realized, she will update beliefs and then choose an act from the menu $F(s_1)$. Thus the way in which she updates will affect the ultimate choice of an act and therefore also the desirability at time 0 of any contingent menu. In this way, the nature of updating is revealed through preference over contingent menus. In particular, because non-Bayesian updating would lead to the "wrong" choice of an act from the menu after s_1 is realized, the agent is led to value commitment at time 0.

Under suitable axioms, we derive a representation for time 0 preference that admits the following interpretation: there are two measures p and q, rather than a single measure as in Anscombe-Aumann and rather than two utility functions as in (1.1). The significance of p is that expected utility relative to p describes preferences under commitment. Therefore, p can be thought of as the counterpart of the Savage or Anscombe-Aumann prior. Updating does not play a role in the ranking of contingent menus that provide commitment because these do not permit any meaningful choice after realization of the signal. Suppose, however, that the agent faces a nonsingleton menu after seeing the signal s_1 and consider the factors influencing her choice of an act from the menu. In analogy with interpretation of the GP functional form (1.1) given above, the second measure q represents the agent's urges in the form of a *retroactively new view* of the world at the interim stage. Her commitment view calls for choosing an act so as to maximize conditional expected utility computed by applying Bayes' Rule to p, but she is tempted to act in accord with her new prior and to maximize expected utility using the Bayesian update of q. In balancing these forces, she behaves *as though* applying Bayes' Rule to a compromise measure p^* that lies between p and q in a suitable sense: each s_1-conditional of p^* is a mixture of the conditionals of p and q, where the mixture weights may vary with the signal s_1. Consequently, interim choice out of the menu given s_1 is based on the compromise posterior $p^*(\cdot \mid s_1)$. If q and p differ, then so also do p^* and p, and updating deviates from application of Bayes' Rule to the commitment prior p.

Note that temptation does *not* refer to whether or not to apply Bayes' Rule. Rather, precisely as in GP, it refers to the temptation to follow one's urges in making choices. The only difference from GP is that here the conflict is between commitment and urges in beliefs (rather than abstract utilities) and it arises after realization of a signal (and not merely with the passage of time).

For an illustration of some of the preceding, consider the following example which adapts GP's motivating example (1.2). The example serves also to illustrate how GP's axiom of Set-Betweenness is adapted below to the present setting. Let

$S_1 = \{s^g, s^b\}$. At time 0 the agent selects a contingent menu of portfolios, that is, a menu for each possible signal. At time 1, after realization of a signal, a portfolio is selected from the menu chosen previously for that state. Finally, there are three possible portfolios - *equity* (consisting exclusively of stocks), a riskless *bond* and *diversified* (*div*), which is a combination of stocks and the bond; each portfolio is an act over S_2 with *bond* being a constant act. Think of s^g (s^b) as constituting good (bad) news about the return to stocks.

Consider the following time 0 ranking of contingent menus:

$$F \equiv \begin{bmatrix} \{equity\} & \text{if } s^g \\ \{div\} & \text{if } s^b \end{bmatrix} \succ \begin{bmatrix} \{equity\} & \text{if } s^g \\ \{bond,\, div\} & \text{if } s^b \end{bmatrix} \succ \begin{bmatrix} \{equity\} & \text{if } s^g \\ \{bond\} & \text{if } s^b \end{bmatrix} \equiv G. \tag{1.3}$$

All contingent menus commit the agent to equity in the event of s^g, while F and G provide perfect commitment also in the bad state. The ranking $F \succ G$ indicates that s^b is only moderately bad news in the sense that it does not justify abandoning stocks entirely. Note that updating is irrelevant to this ranking because there is no interim choice, but it is critical for evaluation of the third contingent menu; denote the latter by $F \cup G$. In particular, interpret the ranking $F \succ F \cup G$ as follows: because the two contingent menus agree given the good signal, the preference between them depends only on what they deliver in s^b. The agent knows her own psyche and anticipates that once the bad signal is realized, she will update in a way that exaggerates the importance (through Bayesian updating of the new prior q). Subsequently, she will be tempted to panic and to leave stocks entirely. This temptation to choose a different portfolio than she would ex ante under commitment, the source of which is her updating behavior, is captured by the strict preference $F \succ F \cup G$. She may anticipate successfully resisting this temptation and choosing *div* from $\{bond,\, div\}$ given s^b, which case is captured by $F \cup G \succ G$.[6] But this choice is contrary to her updated beliefs and feelings of panic and thus requires costly self-control.

An agent with the ranking (1.3) would be willing to pay a positive price to commit to F, say by having her portfolio managed by a suitable investment manager. If committing to F is not possible, it might be possible to avoid seeing the signal. This would limit the agent to choosing contingent menus that do not depend on the signal, but having less information may nevertheless be preferable

[6] Alternatively, she may anticipate succumbing, in which case $F \cup G \sim G$.

if, for example,

$$\begin{bmatrix} \{equity\} & \text{if } s^g \\ \{div\} & \text{if } s^b \end{bmatrix} \succ \begin{bmatrix} \{div\} & \text{if } s^g \\ \{div\} & \text{if } s^b \end{bmatrix}, \text{ but}$$

$$\begin{bmatrix} \{div\} & \text{if } s^g \\ \{div\} & \text{if } s^b \end{bmatrix} \succ \begin{bmatrix} \{equity\} & \text{if } s^g \\ \{bond,\, div\} & \text{if } s^b \end{bmatrix}.$$

One might attempt to interpret (1.3) alternatively in terms of changing tastes or risk aversion. One possibility is that tastes change with the passage of time, independently of which signal is realized. But dependence on the signal is clearly central to the preceding intuition; moreover, the alternative interpretation cannot rationalize a strict preference to ignore information.

Suppose alternatively that tastes change in a state-dependent way. For example, the agent may anticipate becoming more risk averse in response to a bad signal, which might lead to the temptation to choose *bond* from $\{bond,\, div\}$ and hence to the time 0 ranking (1.3). We exclude this interpretation by adopting a suitable axiom of state independence. To illustrate (a special case of) the axiom, let ℓ be a fourth 'security,' thought of as a roulette-wheel whose payoff is independent of the realized state in $S_1 \times S_2$. Suppose that the outcome x is the s^b-conditional certainty equivalent of ℓ in the sense that

$$\begin{bmatrix} \{equity\} & \text{if } s^g \\ \{\ell\} & \text{if } s^b \end{bmatrix} \sim \begin{bmatrix} \{equity\} & \text{if } s^g \\ \{x\} & \text{if } s^b \end{bmatrix}.$$

Then state independence requires that x also be the certainty equivalent conditional on s^g, that is,

$$\begin{bmatrix} \{\ell\} & \text{if } s^g \\ \{equity\} & \text{if } s^b \end{bmatrix} \sim \begin{bmatrix} \{x\} & \text{if } s^g \\ \{equity\} & \text{if } s^b \end{bmatrix}.$$

Where this invariance is accepted, state-dependent risk aversion is excluded, leaving non-Bayesian updating as the only apparent explanation for (1.3).

2. MODEL

2.1. Primitives

The model's primitives include:

- time $t = 0, 1, 2$
- outcome set X (compact metric)

  $\Delta(X)$ denotes the set of lotteries (Borel probability measures) over X

  it is compact metric under the weak convergence topology
- (finite) period state spaces S_1 and S_2 corresponding to the uncertainty resolved at times 1 and 2
- $(\Delta(X))^{S_2}$ is the set of (Anscombe-Aumann) acts over S_2

  the generic act is $f : S_2 \longrightarrow \Delta(X)$
- a closed subset M of $(\Delta(X))^{S_2}$ is called a *menu* (of acts over S_2)

  $\mathcal{M}(S_2)$ is the set of menus

  it is compact metric under the Hausdorff metric[7]
- $F : S_1 \longrightarrow \mathcal{M}(S_2)$ is a *contingent menu*

  $F(s_1)$ is the menu of acts over S_2 from which the agent can choose if s_1 is realized
- $\mathcal{C} = (\mathcal{M}(S_2))^{S_1}$ is the set of all contingent menus
- time 0 preference $\succeq$ is defined on $\mathcal{C}$

The interpretation is that a contingent menu F is chosen ex ante (at time 0) according to $\succeq$. Then at the interim stage $t = 1$, the agent observes the realization of s_1, updates her beliefs about S_2, and finally chooses an act from the menu $F(s_1)$. The state s_2 and hence also the outcome of the chosen act are realized at time 2. Updating and choice behavior at time 1 are anticipated ex ante and underlie the ranking $\succeq$ of contingent menus.

Contingent menus are natural objects of choice.[8] The consequence of a physical action taken at time 0 is that it determines a set of opportunities for further action at time 1, which set depends also on the interim state s_1; that is, the physical action can be identified with a contingent menu. For example, savings at time 0

[7]See [1, Theorem 3.58], for example.

[8]When the menus delivered after realization of a state are menus of alternatives rather than menus of Anscombe-Aumann acts, Kreps [14] proposes contingent menus as the natural objects of choice in a model of unforeseen contingencies; see also Nehring [17].

and the realized state s_1 determine wealth and asset prices, and hence a feasible set of portfolios from which a choice can be made at time 1. A related illustration is provided in (1.3).

Degenerate contingent menus F, where each $F(s_1)$ is a singleton, play a special role. Each such F can be identified with a map $F : S_1 \times S_2 \longrightarrow \Delta(X)$ and thus is an act over $S_1 \times S_2$. The set of such acts is $\mathcal{A} \subset \mathcal{C}$.

Some axioms for preference make use of a notion of nullity that we now introduce. For any act f over S_2 and event $A \subset S_2$, denote by $proj_A f$ the restriction of f to A. Similarly, let

$$proj_A\, M \;\equiv \{proj_A\, f : \; f \in M\},$$

for any $A \subset S_2$ and menu M of acts over S_2; when $A = \{s_2\}$, write simply $proj_{s_2}$. Below we refer to $proj_A\, M$ as the *projection* of M on A.

For any s_1 in S_1 and event $E \subset S_2$, say that (s_1, E) is *null* if $F' \sim F$ for all F' and F satisfying both

$$F'(s_1') \;=\; F(s_1') \quad \text{for all } s_1' \neq s_1 \tag{2.1}$$

and

$$proj_{S_2 \setminus E}\, F'(s_1) \;=\; proj_{S_2 \setminus E}\, F(s_1)\,. \tag{2.2}$$

In words, (s_1, E) is null if any two contingent menus that "differ only on $\{s_1\} \times E$" are indifferent. In particular, (s_1, S_2) is null if $F' \sim F$ whenever (2.1) is satisfied. When $E = \{s_2\}$ is a singleton, refer to nullity of (s_1, s_2) rather than of $(s_1, \{s_2\})$.

2.2. Axioms

Turn to axioms for $\succeq$. The first two need no explanation.

Axiom 1 (Order). *$\succeq$ is complete and transitive.*

Axiom 2 (Continuity). *The sets $\{F \in \mathcal{C} : F \succeq G\}$ and $\{F \in \mathcal{C} : F \preceq G\}$ are closed.*

The set $(\Delta(X))^{S_2}$ of Anscombe-Aumann acts over S_2 is a mixture space. Further, any two menus of such acts, M and N, can be mixed according to

$$\lambda M + (1-\lambda) N \;=\; \{\lambda f + (1-\lambda) g : \; f \in M,\, g \in N\}\,.$$

Finally, for any two contingent menus F and G, define the mixture statewise by

$$(\lambda F + (1-\lambda) G)\ (s_1)\ = \lambda\, F\,(s_1) + (1-\lambda)\, G\ (s_1)\,, \quad s_1 \in S_1.$$

We can now state the Independence Axiom for our setting.

Axiom 3 (Independence). *For every* $0 < \lambda \leq 1$, $F \succeq G$ *iff* $\lambda F + (1-\lambda) F' \succeq \lambda G + (1-\lambda) F'$.

A first stab at intuition for Independence is similar to that familiar from the Anscombe-Aumann model and also to that offered in [6, 10] for their versions of the axiom. For completeness, we describe it briefly. The mixture $\lambda F + (1-\lambda) F'$ is the contingent menu that delivers the set of acts $\lambda F(s_1) + (1-\lambda) F'(s_1)$ in state s_1. Consider instead the lottery over $\mathcal{C}$, denoted $\lambda \circ F + (1-\lambda) \circ F'$, that delivers F with probability λ and F' with probability $(1-\lambda)$. Supposing that the agent can rank such lotteries, then the familiar intuition for the usual form of Independence suggests that $F \succeq G$ iff $\lambda \circ F + (1-\lambda) \circ F' \succeq \lambda \circ G + (1-\lambda) \circ F'$. Thus the intuition for our version of Independence is complete if we can justify indifference between $\lambda \circ F + (1-\lambda) \circ F'$ and $\lambda F + (1-\lambda) F'$. The difference between them is that under the former, randomization is completed immediately, at $t = 0$, while for the latter, the timing is such that s_1 is realized, (beliefs are updated), and then the agent chooses an act from the convex combination $\lambda F(s_1) + (1-\lambda) F'(s_1)$ of menus of acts. The latter corresponds also to the randomization with weight λ occurring after the interim choice of an act. Thus the desired indifference amounts to indifference to the timing of resolution of uncertainty. (Dekel, Lipman and Rustichini (pp. 905-6) provide a normative justification for indifference to timing that can be adapted to the present setting.)

However, there is more implicit in Independence. Consider the lottery $\lambda \circ F + (1-\lambda) \circ F'$. After the randomization is completed, the agent updates her beliefs over $S_1 \times S_2$. Though the randomization is objectively independent of events in $S_1 \times S_2$, given that she changes her view of the world after making an observation, the agent might change her beliefs over $S_1 \times S_2$. As a result, she might prefer $\lambda \circ G + (1-\lambda) \circ F'$ to $\lambda \circ F + (1-\lambda) \circ F'$ even while preferring F to G. Thus intuition for Independence assumes that, consistent with our agent not being one who makes mistakes, she recognizes the *objective* fact that randomization is unrelated to the state space. At the same time, she may, according to our model, view the events $E_1 \subset S_1$ and $E_2 \subset S_2$ as *subjectively* independent according to her initial (commitment) prior and yet change her beliefs about E_2 after seeing E_1.

To rule out trivial cases, adopt:[9]

Axiom 4 (Nondegeneracy). *There exist x, y in X for which $x \succ y$.*

Any state s_1 such that (s_1, S_2) is null could simply be dropped. Thus there is no loss of generality in assuming that no such states exist.

Axiom 5 (S_1-Non-nullity). *For every s_1, (s_1, S_2) is non-null.*

At this point we depart from Anscombe-Aumann. While their model can be viewed as (implicitly) imposing a form of strategic rationality (see the discussion following Theorem 3.1), in order to permit temptation and self-control we adopt a counterpart of Gul and Pesendorfer's Set-Betweenness axiom. To state the axiom, define the union $F \cup G$ statewise, that is,

$$(F \cup G)(s_1) = F(s_1) \cup G(s_1).$$

Axiom 6 (Set-Betweenness). *For all states s_1 and all menus F and G such that $G(s_1') = F(s_1')$ for all $s_1' \neq s_1$,*

$$F \succeq G \implies F \succeq F \cup G \succeq G. \tag{2.3}$$

Because F and G are identical in all states $s_1' \neq s_1$, $F \succeq G$ means that given s_1 at the interim stage, the agent would rather have $F(s_1)$ than $G(s_1)$ from which to choose after updating. Conditional preference over menus at any s_1 is derived from the subsequent choice of acts that is anticipated to follow immediately, as in the GP model. Thus the identical motivation offered by GP (p. 1408) applies here. In particular, the hypothesis that temptation cannot increase utility and that the utility cost of temptation depends only on the most tempting alternative, leads to the agent's conditional preference at the interim stage for $F(s_1)$ over $F(s_1) \cup G(s_1)$ and preference for the latter over $G(s_1)$. But F, $F \cup G$ and G coincide in all states $s_1' \neq s_1$ and thus, from the ex ante perspective, the desired ranking of $F \cup G$ follows. The portfolio choice example (1.3) illustrates the preceding.

[9] $x \in X$ is identified with the contingent menu that, in every state s_1, yields the (singleton menu comprised of the) lottery yielding x with probability 1.

For perspective, consider the stronger axiom that would impose (2.3) for all contingent menus and not just for those that differ only in one state s_1. It is easily seen that this stronger axiom is not intuitive. For example, suppose that

$$F \equiv \begin{bmatrix} \{f\} & \text{if } s_1 \\ \{f'\} & \text{if } s_1' \end{bmatrix} \succ \begin{bmatrix} \{g\} & \text{if } s_1 \\ \{g'\} & \text{if } s_1' \end{bmatrix} \equiv G,$$

where f and g' are very attractive acts over S_2 while f' and g are less attractive but suitably tempting. Suppose that

$$\{f\} \succ_{s_1} \{f, g\} \sim_{s_1} \{g\} \text{ and}$$
$$\{g'\} \succ_{s_1'} \{f', g'\} \sim_{s_1'} \{f'\},$$

where $\succeq_{s_1}$ and $\succeq_{s_1'}$ denote preference at the interim stage given realization of s_1 and s_1' respectively. In particular, g is so tempting given s_1 that it would be chosen out of $\{f, g\}$, and f' is so tempting given s_1' that it would be chosen out of $\{f', g'\}$. Therefore, $F \cup G$ would lead ultimately to the choice of g given s_1 and f' given s_1', the worst of both worlds, which suggests the ranking $G \succ F \cup G$.

The next axiom is the principal way in which temptation is connected to changing beliefs. At the functional form level, the axiom is important in tracing the difference between the counterparts of the two GP functions u and v in (1.1) to differences in beliefs rather than to differences in risk attitudes or utilities over final outcomes. Its statement requires additional notation. For any act f over S_2, lottery ℓ and state s_2, denote by $\ell s_2 f$ the act over S_2 that assigns ℓ if the realized state is s_2 and $f(s_2')$ if the realized state is $s_2' \neq s_2$. Similarly, for any menu $M \in \mathcal{M}(S_2)$ and menu of lotteries $L \subset \Delta(X)$,

$$Ls_2M \equiv \{\ell s_2 f : \ell \in L, \, f \in M\}. \tag{2.4}$$

A distinguishing feature of such a menu of acts over S_2 is that it "equals the Cartesian product of its projections on $\{s_2\}$ and $S_2 \backslash \{s_2\}$."

To illustrate this notation, consider the example in the introduction and let $S_2 = \{s_2, s_2'\}$. Recall that *div* is an act over S_2; denote by $div(s_2)$ the payoff to the diversified portfolio in state s_2, and so on. The *bond* is a constant act, that is, a lottery. Let ℓ be any other lottery, $L = \{bond, \ell\}$ and $M = \{equity, div\}$. Then $Ls_2'M$ consists of the following 4 acts:[10]

$$\begin{bmatrix} div(s_2) \\ bond \end{bmatrix}, \begin{bmatrix} div(s_2) \\ \ell \end{bmatrix}, \begin{bmatrix} equity(s_2) \\ bond \end{bmatrix}, \begin{bmatrix} equity(s_2) \\ \ell \end{bmatrix}. \tag{2.5}$$

[10] Acts are 2-vectors where the components give payoffs in states s_2 and s_2' respectively.

The significance of the noted Cartesian product structure is as follows. Consider the agent after s_1 is realized and facing the menu Ls_2M of acts over S_2. In evaluating the menu, she anticipates updating to incorporate the observed signal s_1 and then choosing an act from the menu Ls_2M. Her payoff is then determined by the chosen act and the realized state in S_2. Though the choice of an act is made before learning whether or not s_2 is true, menus of the above form permit the full range of contingent choices that would be possible if choice could be made 'ex post' after learning if s_2 is true, (as is evident in the portfolio example). Thus we can equally well think of choice as being made ex post and of the agent as having the following perspective: if s_2 is realized, then I will choose a lottery from L, and if s_2 is not realized, then I will choose an act from M. Similarly when evaluating $L's_2M$. Thus when comparing $L's_2M$ and Ls_2M, the usual intuition for separability across disjoint events suggests that the comparison reduces to the question "given state s_2, would I rather choose a lottery from L' or from L?" This perspective is exploited in several instances below.

Denote by (F_{-s_1}, Ls_2M) the contingent menu that delivers $F(s_1')$ is $s_1' \neq s_1$ and Ls_2M otherwise.

Axiom 7 (Restricted Strategic Rationality (RSR)). *There exists a non-null state (s_1, s_2), such that for all contingent menus F, menus $M \in \mathcal{M}(S_2)$ and menus of lotteries L', $L \subset \Delta(X)$,*

$$(F_{-s_1}, L's_2M) \succ (F_{-s_1}, Ls_2M) \implies$$

$$(F_{-s_1}, L's_2M) \sim (F_{-s_1}, (L' \cup L)\, s_2M).$$

Given the axiom State Independence specified below, RSR could be equivalently stated as "for all non-null (s_1, s_2) ..." in place of "there exists non-null (s_1, s_2) ..." Note also that the invariance indicated in RSR is satisfied vacuously if (s_1, s_2) is null.

To interpret the axiom, consider the ranking of $G' = (F_{-s_1}, L's_2M)$ and $G = (F_{-s_1}, Ls_2M)$. By the intuition underlying the Sure-Thing-Principle, the agent evaluates them ex ante by considering how she would rank the menus $G'(s_1')$ and $G(s_1')$ upon realization of any state s_1'. The issue is whether updating is relevant for making any of these interim rankings. The answer is clearly no for any $s_1' \neq s_1$, because then she faces the same menus given either G' or G. The claim is that updating is irrelevant also in state s_1, where the comparison is between $L's_2M$ and Ls_2M. We have just seen that an agent facing these two menus after realization of s_1, or anticipating this situation ex ante, sees two menus that "differ only in the single state s_2." Evidently, such a comparison does not involve trade-offs across states in S_2 and hence does not depend on beliefs or likelihoods, which in turn means that there is no need to update beliefs about S_2 in response to the received signal s_1. Because temptations in choice arise only with a change of beliefs, a form of strategic rationality should be valid for such comparisons.

As in Anscombe-Aumann, a form of state independence is needed.

Axiom 8 (State Independence). *For all $s = (s_1, s_2)$ and non-null $s' = (s_1', s_2')$, contingent menus F, menus M of acts over S_2 and menus L' and L of lotteries,*

$$\left(F_{-s_1'}, L' s_2' M\right) \succeq \left(F_{-s_1'}, L s_2' M\right) \implies \left(F_{-s_1}, L' s_2 M\right) \succeq \left(F_{-s_1}, L s_2 M\right).$$

The ranking on the left indicates the preference to choose a lottery from L' rather than from L given the state (s_1', s_2'). The indicated ranking could reflect nullity of (s_1', s_2'), but this is excluded by assumption. We saw in the discussion of RSR that the ranking does not reflect beliefs, and RSR implies that temptation is not relevant. Thus it is presumably taste, or risk aversion, that underlies the ranking. But if risk aversion is independent of the state, then the corresponding ranking should prevail also for (s_1, s_2).

Axiom 9 (Absolute Continuity). *For all (s_1, s_2) and contingent menus F, if*

$$\left(F_{-s_1}, L' s_2 F(s_1)\right) \sim \left(F_{-s_1}, L s_2 F(s_1)\right)$$

for all *menus of lotteries L' and L, then*

$$F' \equiv \left(F_{-s_1}, L s_2 F(s_1)\right) \sim F, \text{ where } L = proj_{s_2} F(s_1). \tag{2.6}$$

Note that, given the hypothesis, the conclusion (2.6) is equivalent to asserting that (s_1, s_2) is null.

To interpret the axiom, compare F' and F in (2.6). These contingent menus have much in common - they agree for all $s_1' \neq s_1$ and, for state s_1,

$$proj_A\, F'(s_1) = proj_A\, F(s_1)\,, \text{ for } A = \{s_2\},\ S_2 \backslash \{s_2\}.$$

But this is not enough to make indifference intuitive. One formal difference between F' and F is that only $F'(s_1)$ has a "Cartesian product structure" such as described prior to the statement of RSR. As explained there, the significance of this structure for a menu has to do with the time at which choice of an act from the menu can be thought to be made. For the typical menu $F(s_1)$, choice of an act from the menu must be made before learning whether s_2 is true. However, $F'(s_1) = Ls_2F(s_1)$ is sufficiently rich that it can replicate any contingent choice, and thus it is as though choice of an act from $F'(s_1)$ can be made after learning whether or not s_2 is true. Thus F' and F provide substantially different menus given s_1, and $F' \sim F$ is not intuitive.

Suppose, however, that the hypothesis of the axiom is satisfied. Then because the agent does not care which menu of lotteries is available if (s_1, s_2) is realized, she must believe ex ante that (s_1, s_2) cannot happen. Suppose that she anticipates that she will continue to view s_2 as impossible also if s_1 is realized (roughly, that conditional beliefs are absolutely continuous with respect to ex ante beliefs).[11] Then she anticipates that given s_1 and facing $F(s_1)$, she will believe that s_2 is impossible and thus will find herself 'as if' she can choose from $F(s_1)$ after learning if s_2 is true. But this is precisely the situation she would anticipate if she expected that state s_1 would lead to the menu $Ls_2F(s_1)$, where $L = proj_{s_2}F(s_1)$. Conclude that she will be indifferent between F and $(F_{-s_1}, Ls_2F'(s_1))$ ex ante.

While the preceding axioms lead to a sharp representation, below we sometimes assume also the following axioms. For any act f over S_2, denote by $\overline{\{f\}}$ the contingent menu that commits the agent to f in every state s_1. Evidently, the evaluation of any such prospect reflects marginal beliefs about S_2 held at time 0, that is, the agent's prior on S_2.

Axiom 10 (Prior-Bias). *For each s_1, for any contingent menu F and for any f and g, acts over S_2, if*

$$(F_{-s_1}, \{f\}) \succ (F_{-s_1}, \{g\}) \text{ and } \overline{\{f\}} \sim \overline{\{g\}},$$

then

$$(F_{-s_1}, \{f\}) \sim (F_{-s_1}, \{f, g\}).$$

The hypothesis indicates that the agent strictly prefers to commit to f rather than to g for the eventuality where s_1 is realized, and that she is indifferent between them according to her prior beliefs on S_2. Under these conditions, she is not tempted by g conditionally on s_1. That the absence of temptation conditionally on s_1 depends not only on how f and g are ranked conditionally (as in the first

[11]Note that by S_1-Non-nullity, realization of s_1 would not be a total surprise.

part of the hypothesis) but also on how attractive they were prior to realization of s_1, indicates excessive influence of prior beliefs at the updating stage.

Prior-Bias permits prior beliefs to unduly influence updating but does not specify the direction of such influence. Put another way, what happens if $\overline{\{f\}}$ and $\overline{\{g\}}$ are not indifferent? We consider two alternative strengthenings of the axiom that provide different answers.

Axiom 11 (Positive Prior-Bias). *For each s_1, for any contingent menu F and for any f and g, acts over S_2, if*

$$(F_{-s_1}, \{f\}) \succ (F_{-s_1}, \{g\}) \text{ and } \overline{\{f\}} \succeq \overline{\{g\}},$$

then

$$(F_{-s_1}, \{f\}) \sim (F_{-s_1}, \{f, g\}).$$

According to this axiom, g is tempting conditionally on s_1 only if it was more attractive according to prior beliefs on S_2. Consequently, temptation can arise when conditional and ex ante rankings disagree strictly.

An alternative strengthening of Prior-Bias is:

Axiom 12 (Negative Prior-Bias). *For each s_1, for any contingent menu F and for any f and g, acts over S_2, if*

$$(F_{-s_1}, \{f\}) \succ (F_{-s_1}, \{g\}) \text{ and } \overline{\{f\}} \preceq \overline{\{g\}},$$

then

$$(F_{-s_1}, \{f\}) \sim (F_{-s_1}, \{f, g\}).$$

The hypothesis states that while g was (weakly) preferred according to prior beliefs on S_2, the signal s_1 reversed the ranking in favor of f. This agent is greatly influenced by, or overreacts to, such a strong signal (and she knows this about herself ex ante). Thus she is not at all tempted by g after seeing s_1 (or when anticipating its realization ex ante). Of course, the axiom leaves open the possibility of temptation when $\overline{\{f\}} \succ \overline{\{g\}}$.

3. UTILITY

3.1. Representation Result

Define the utility function $\mathcal{U}$ on $\mathcal{C}$ in two stages. First, evaluate F via the (state-dependent) expected utility form

$$\mathcal{U}(F) = \int_{S_1} \mathcal{U}(F(s_1); s_1)\, dp_1, \quad F \in \mathcal{C}, \tag{3.1}$$

where p_1 is a probability measure on S_1 and each $\mathcal{U}(\cdot; s_1)$ is a utility function on the collection of menus of acts over S_2. Its specification is the heart of the model. Think of $\mathcal{U}(\cdot; s_1)$ as the utility of menus after s_1 is realized but *before updating* is performed.

The GP utility functional form (1.1) suggests the form

$$\mathcal{U}(F(s_1); s_1) = \max_{f \in F(s_1)} \{U(f; s_1) + V(f; s_1)\} - \max_{f' \in F(s_1)} V(f'; s_1), \tag{3.2}$$

for suitable functions $U(\cdot; s_1)$ and $V(\cdot; s_1)$. The particular specification that we adopt is

$$\mathcal{U}(F(s_1); s_1) = \max_{f \in F(s_1)} \left\{ \int_{S_2} u(f)\, dp(\cdot \mid s_1) + \alpha(s_1) \int_{S_2} u(f)\, dq(\cdot \mid s_1) \right\} \tag{3.3}$$

$$- \max_{f' \in F(s_1)} \alpha(s_1) \int_{S_2} u(f')\, dq(\cdot \mid s_1),$$

where components of the functional from satisfy the *regularity conditions*:

Reg1 $u : \Delta(X) \longrightarrow \mathbb{R}^1$ is mixture linear, continuous and nonconstant.

Reg2 Each $p(\cdot \mid s_1)$ and $q(\cdot \mid s_1)$ is a probability measure on S_2, $q(\cdot \mid s_1)$ is absolutely continuous with respect to $p(\cdot \mid s_1)$.

Reg3 $\alpha : S_1 \longrightarrow [0, \infty)$.

Assume in addition that the measure p_1 appearing in (3.1) satisfies:

Reg4 p_1 has full support on S_1.

Utility is defined by (3.1), (3.3) and the regularity conditions.

Let p be the measure on $S_1 \times S_2$ generated by p_1 and the conditionals $\{p(\cdot \mid s_1) : s_1 \in S_1\}$. It is convenient also to define the measure q generated by p_1 and the conditionals $\{q(\cdot \mid s_1) : s_1 \in S_1\}$. Then p_1 is the S_1-marginal of p, $p(\cdot \mid s_1)$ is the Bayesian conditional of p, and similarly for q. Further, q and p have identical S_1-marginals and $q << p$.[12]

The representation admits the interpretation outlined in the introduction. For prospects that offer commitment, that is, if $F \in \mathcal{A}$, utility simplifies to

$$\mathcal{U}(F) = \int_{S_1 \times S_2} u(F)\ dp, \text{ for } F \in \mathcal{A}. \tag{3.4}$$

Thus p is the *commitment prior*, and is the counterpart of the usual prior.

For a general contingent menu F, $\mathcal{U}(F(s_1); s_1) =$

$$\max_{f \in F(s_1)} \left\{ \int_{S_2} u(f)\, dp(\cdot \mid s_1) + \alpha(s_1) \left[\int_{S_2} u(f)\, dq(\cdot \mid s_1) - \max_{f' \in F(s_1)} \int_{S_2} u(f')\, dq(\cdot \mid s_1) \right] \right\}.$$

Thus choice of an act from the menu $F(s_1)$ is based only in part on the Bayesian update of p. Another factor is that the agent retroactively adopts the revised prior q and thus is tempted to choose an act that solves $\max_{f' \in F(s_1)} \int_{S_2} u(f')\, dq(\cdot \mid s_1)$. To the extent that she resists this temptation and chooses another act f, she incurs the (utility) self-control cost

$$\alpha(s_1) \left[\max_{f' \in F(s_1)} \int_{S_2} u(f')\, dq(\cdot \mid s_1) - \int_{S_2} u(f)\, dq(\cdot \mid s_1) \right];$$

$\alpha(s_1)$ parametrizes the cost of self-control in state s_1 (see further discussion in the next section). Compromise between the commitment perspective and the cost of self-control leads to choice from the menu according to

$$max_{f \in F(s_1)} \left\{ \int_{S_2} u(f)\ dp(\cdot \mid s_1) + \alpha(s_1) \int_{S_2} u(f)\ dq(\cdot \mid s_1) \right\}. \tag{3.5}$$

Finally, define p^* on $S_1 \times S_2$ by

$$p^*(s_1, s_2) = \tfrac{p(s_2|s_1) + \alpha(s_1) q(s_2|s_1)}{1+\alpha(s_1)} p_1(s_1). \tag{3.6}$$

[12] $q << p$ means that q is absolutely continuous with respect to p. Below we denote by p_2 the marginal of p on S_2.

Then choice from the menu is made as though maximizing expected utility using the Bayesian update of p^*, where p^* can be thought of as a *compromise prior* and

$$p^*(\cdot \mid s_1) = \tfrac{p(\cdot|s_1)+\alpha(s_1)q(\cdot|s_1)}{1+\alpha(s_1)}. \tag{3.7}$$

The latter differs from the Bayesian update of p to the extent that the conditionals $q(\cdot \mid s_1)$ and $p(\cdot \mid s_1)$ differ.

The central result of the paper is the following axiomatization of utility over contingent menus:

Theorem 3.1. *(a) $\succeq$ satisfies Order, Continuity, Independence, Nondegeneracy, S_1-Non-nullity, Set-Betweenness, State Independence, RSR and Absolute Continuity if and only if it admits a representation of the form (3.1)-(3.3), including the regularity conditions Reg1-Reg4.*

(b) Moreover, $\succeq$ satisfies also Prior-Bias (Positive Prior-Bias or Negative Prior-Bias, respectively) if and only if it admits a representation as above where in addition: for each s_1, either $\alpha(s_1) = 0$, or $\alpha(s_1) > 0$ and

$$q(\cdot \mid s_1) = (1 - \lambda(s_1))\, p(\cdot \mid s_1) + \lambda(s_1)\, p_2(\cdot), \tag{3.8}$$

where $\lambda(s_1) \leq 1$ ($0 < \lambda(s_1) \leq 1$, $\lambda(s_1) \leq 0$, respectively).

Focus on part (a) and on the representation (3.1)-(3.3). For perspective, consider an alternative functional form for utility which satisfies (3.1)-(3.2) with

$$U(f; s_1) = \int_{S_2} u(f)\, dp(\cdot \mid s_1),$$

but where the specification of temptation utility $V(\cdot; s_1)$ is modified so that $\mathcal{U}(F(s_1); s_1) =$

$$\left\{ \max_{f \in F(s_1)} \int_{S_2} [u(f) + v(f)]\ dp(\cdot \mid s_1) \right\} - \max_{f' \in F(s_1)} \int_{S_2} v(f')\ dp(\cdot \mid s_1).$$

Here there is a single probability measure p but two utility indices u and v. This model satisfies all axioms but RSR. Violation of RSR indicates that temptation is an issue even where menus are such that beliefs and hence updating are irrelevant to their evaluation. This suggests that the functional form captures temptation that arises from changes in taste that occur with the passage of time, along the lines of Strotz [20], rather than because of the nature of updating. Appendix B describes axiomatic foundations for this model (augmented by suitable regularity conditions) in order to permit a clearer comparison with, and perspective for, our model.

The relation of the theorem to the (dynamic) Anscombe-Aumann model merits emphasis. The latter is obtained if one strengthens Set-Betweenness to strategic rationality, that is, if one requires that $F \succeq G \implies F \sim G$ whenever F and G differ only in one state s_1.

A final comment concerns generality of the model (3.1)-(3.3). Because it imposes little structure on the relation between p and q (or equivalently between p and the compromise prior p^*), the model can accommodate a range of updating biases (see Section 3.3). On the other hand, it may be "too general" in that it permits beliefs to change ($q \neq p$) even when S_1 is a singleton and thus when there is no real signal.[13] Partly for this reason we have specialized the model by adding Prior-Bias. The latter implies that, for any given s_1, *either* $\alpha(s_1) = 0$ and updating in response to s_1 amounts to Bayesian updating of p, *or* $q(\cdot \mid s_1)$ has the form (3.8). In particular, $q(\cdot \mid s_1) = p(\cdot \mid s_1)$ and updating is standard (albeit trivial) if S_1 is a singleton, or more generally, if p is a product measure.

The functional form (3.8) is readily interpreted. Suppose, for example, that $0 \leq \lambda(s_1) \leq 1$ corresponding to Positive Prior-Bias. Then $q(\cdot \mid s_1)$ is a mixture of $p(\cdot \mid s_1)$ and prior beliefs p_2. Because $p(\cdot \mid s_1)$ embodies "the correct" combination of prior beliefs and responsiveness to data, and because p_2 gives no weight to data, the updating implied by (3.8) gives "too much" weight to prior beliefs and "too little" to observation. (Note that the same is true for the compromise posterior $\frac{p(\cdot|s_1)+\alpha q(\cdot|s_1)}{1+\alpha}$ from (3.7).) Prior beliefs exert undue influence, relative to Bayesian updating, also if $\lambda(s_1) < 0$. To interpret this case, fix a state s_2 and rewrite (3.8) in the form

$$q(s_2 \mid s_1) = p(s_2 \mid s_1) - \lambda(s_1)(p(s_2 \mid s_1) - p_2(s_2)).$$

If $p(s_2 \mid s_1) - p_2(s_2) > 0$, then s_1 is a strong positive signal for s_2. In this case the agent overreacts to such positive signals to a degree described by $-\lambda(s_1)$. Prior beliefs have an undue (negative) influence in that they are already taken into account to a proper degree in $p(s_2 \mid s_1)$.

[13] This case is ruled out if $S_1 = S_2 = S$, which is a common specification (repeated experiments). See Section 3.3 for examples.

3.2. Uniqueness and Further Interpretation

Say that (u, p, q, α) *represents* $\succeq$ if it satisfies the conditions part (a) of the theorem. Next we describe the uniqueness properties of such representing tuples under an additional assumption. To simplify statement of the latter, define the conditional order $\succeq_{s_1}$ on $\mathcal{M}(S_2)$ by

$$M' \succeq_{s_1} M \quad \text{iff } \exists F \text{ such that } \left(F_{-s_1}, M'\right) \succeq \left(F_{-s_1}, M\right).$$

Given the other axioms, "$\exists F$" is equivalent to "$\forall F$" and $\succeq_{s_1}$ is represented by $\mathcal{U}(\cdot; s_1)$ defined in (3.3). Though $\succeq_{s_1}$ is defined as an ex ante ranking, we interpret it also as the preference that would prevail at the interim stage after realization of s_1 *but before updating*.

The following (elementary) lemma describes several equivalent statements of the needed additional assumption.

Lemma 3.2. *Let $\succeq$ satisfy the axioms in part (a) of the theorem and fix s_1 in S_1. Then the following statements are equivalent:*

(a) There exist menus M' and M such that $M' \succeq_{s_1} M$ and $M' \nsim_{s_1} M' \cup M$.

(b) There exist f and g, Anscombe-Aumann acts over S_2, such that

$$\{f\} \succ_{s_1} \{f, g\} \succ_{s_1} \{g\}\ . \tag{3.9}$$

(c) There exists an Anscombe-Aumann act f over S_2 and a lottery ℓ, such that

$$\{f\} \succ_{s_1} \{f, \ell\} \succ_{s_1} \{\ell\}\ . \tag{3.10}$$

(d) There exists a representing tuple (u, p, q, α) such that

$$\alpha(s_1) \neq 0 \quad \text{and } q(\cdot \mid s_1) \neq p(\cdot \mid s_1)\,. \tag{3.11}$$

Part (a) states that $\succeq_{s_1}$ violates strategic rationality. Thus it excludes the case where conditional utility $\mathcal{U}(\cdot; s_1)$ as defined in (3.3) takes the form

$$\mathcal{U}(M; s_1) = \max_{f \in M} \int_{S_2} u(f)\, dp(\cdot \mid s_1),$$

for any menu M of acts over S_2, precisely as in the standard model with Bayesian updating. In that sense, each of the conditions in the Lemma amounts to the assumption of non-Bayesian updating given s_1. In the terminology of [10, p.

1413], (b) states that the agent *has self-control at* $\{f, g\}$ conditionally on s_1. Part (c) asserts the existence of such self-control where g is a lottery (constant act). Finally, (d) provides the corresponding restrictions on the representing functional form.[14]

Corollary 3.3. *Let* (u, p, q, α) *represent* $\succeq$*. Then* (u', p', q', α') *also represents* $\succeq$ *if and only if: (i) there exists* $(a, b) \in \mathbb{R}^1_{++} \times \mathbb{R}^1$ *such that*

$$u' = au + b \text{ and } \; p' = p;$$

and *(ii) for every* s_1*, either*

$$\alpha'(s_1)\left(q'(\cdot \mid s_1) - p'(\cdot \mid s_1)\right) = 0 = \alpha(s_1)\left(q(\cdot \mid s_1) - p(\cdot \mid s_1)\right), \text{ or} \tag{3.12}$$

$$\alpha'(s_1) = \alpha(s_1) \quad \text{and } q'(\cdot \mid s_1) = q(\cdot \mid s_1). \tag{3.13}$$

The uniqueness properties of (u, p) are straightforward and expected. For (q, α), the relevant uniqueness property depends on s_1. One possibility is (3.12) which states that both representations violate condition (3.11). In that case, interim choice behavior is based on the Bayesian update of p. Otherwise, the strong uniqueness statement (3.13) is valid for s_1.

If conditions of the Lemma are satisfied for every s_1, then (u', p', q', α') and (u, p, q, α) both represent $\succeq$ if and only if

$$u' = au + b \text{ and } (p', q', \alpha') = (p, q, \alpha),$$

for some $a > 0$ and $b \in \mathbb{R}^1$. Then q and $\alpha(\cdot)$, in addition to p, are unique and hence meaningful components of the functional form. It makes sense then to consider their behavioral meaning. For p, we have already observed that it is the prior that guides choice under commitment. The meaning of α can be described explicitly under conditions of the Lemma as we now show.[15]

Let x^{**} and x^* be best and worst alternatives under commitment, that is, such that

$$\{x^{**}\} \succeq \{x\} \succeq \{x^*\} \quad \text{for all } x \text{ in } X.$$

(They exist by Continuity and compactness of X.) Then also

$$\{x^*\} \succeq_{s_1} M \succeq_{s_1} \{x^{**}\}$$

[14] Prove that (d) implies (c): From (d), $\int u(f)\, dq(\cdot \mid s_1) < \int u(f)\, dp(\cdot \mid s_1)$ for some f. Choose ℓ such that $u(\ell) < \frac{\int u(f)\, dp(\cdot|s_1) + \alpha(s_1) \int u(f)\, dq(\cdot|s_1)}{(1+\alpha(s_1))}$.

[15] The explication of q is straightforward but omitted for brevity.

for all states s_1 and menus M. Normalizing u so that

$$u(x^*) = 0 \text{ and } u(x^{**}) = 1,$$

then, as in vNM theory, utilities are directly observable as 'mixture weights'. That is, because each $\mathcal{U}(\cdot; s_1)$ is mixture linear, $\mathcal{U}(M; s_1)$ is the unique weight m such that

$$m\{x^{**}\} + (1-m)\{x^*\} \sim_{s_1} M. \tag{3.14}$$

Similarly for the special case $\mathcal{U}(\{\ell\}; s_1) = u(\ell)$.

This permits isolation of the behavioral meaning of $\alpha(\cdot)$, as described in the following corollary.

Corollary 3.4. *Suppose that $\succeq_{s_1}$ satisfies conditions of the Lemma and let there be self-control at both $\{f, \ell\}$ and $\{f, \ell'\}$ as in (3.10), where $\{\ell'\} \nsim_{s_1} \{\ell\}$. Then*

$$\alpha(s_1) = \frac{\mathcal{U}(\{f, \ell\}; s_1) - \mathcal{U}(\{f, \ell'\}; s_1)}{u(\ell') - u(\ell)}. \tag{3.15}$$

For lotteries ℓ and ℓ' as in the statement, compute that

$$\mathcal{U}(\{f\}; s_1) - \mathcal{U}(\{f, \ell\}; s_1) = \alpha(s_1)\left[u(\ell) - \int_{S_2} u(f)\; dq(\cdot \mid s_1)\right],$$

and similarly for ℓ'. Expression (3.15) follows. It is important to note that each utility level appearing on the right side of (3.15) is observable from behavior using (3.14). Thus we have a *closed-form and behaviorally meaningful* expression for $\alpha(s_1)$. Because each mixture weight appearing in (3.14) is unit-free, so is the expression given for $\alpha(s_1)$.

For further interpretation of $\alpha(s_1)$, GP's Theorem 9, translated to our setting, yields that $\succeq_{s_1}$ (satisfying conditions in the Lemma) exhibits less self-control the larger is $\alpha(s_1)$.[16] Following GP, $\succeq_{s_1}$ exhibits less self-control than $\succeq'_{s_1}$ if, for all menus M and N of acts over S_2, $M \succ_{s_1} M \cup N \succ_{s_1} N$ implies the same ranking in terms of $\succ'_{s_1}$. In addition to this interpretation in terms of comparative self-control, the expression (3.15) permits interpretation of $\alpha(s_1)$ as an absolute measure of self-control. If there is self-control at $\{f, \ell\}$, as in (3.10), then $\mathcal{U}(\{f\}; s_1) - \mathcal{U}(\{f, \ell\}; s_1)$ is the utility cost of having ℓ available and exerting

[16] Theorem 9 seems misstated - the correct statement should fix the parameter γ to equal 1. In their sequel [11, Theorem 5], the authors refer to this corrected version of Theorem 9.

self-control in order to choose f, where utility is measured in probabilities as in (3.14). Thus $\alpha(s_1)$ gives the rate at which this self-control cost increases as ℓ improves in the sense measured by $u(\ell)$. In that sense, $\alpha(s_1)$ is the *marginal cost of self-control* in state s_1.

3.3. Examples of Updating Biases

This section demonstrates the richness of our model by showing how it can produce, through suitable specifications for p, q and α, a variety of updating biases, including some that are analogous to biases discussed by psychologists in the context of objective probabilities. Our claim here is not that we can accommodate all or many of these with a single specification, though future research will explore that possibility. For now we content ourselves with suggesting the potential of our model to provide a unifying and choice-theoretic framework.

Underreaction and Overreaction: Our discussion of the axioms Positive Prior-Bias and Negative Prior-Bias and of the functional form (3.8) suggested an agent who underreacts (overreacts) to observations. This interpretation is clearer if we assume also that $\lambda(\cdot)$ and $\alpha(\cdot)$ are constant. Then choice after realization of s_1 is based on the Bayesian update of the compromise prior, and hence, by (3.7), on the conditional measure

$$p^*(\cdot \mid s_1) = \left(1 - \tfrac{\alpha\lambda}{1+\alpha}\right) p(\cdot \mid s_1) + \tfrac{\alpha\lambda}{1+\alpha} p_2(\cdot). \tag{3.16}$$

Evidently, $p^*(\cdot \mid s_1)$ is less sensitive to the signal s_1 than is $p(\cdot \mid s_1)$ if $\lambda > 0$ and more sensitive if $\lambda < 0$. In particular, $\lambda < 0$ can capture the temptation to panic in the face of the bad signal s^b as discussed in the introductory portfolio choice example. The larger is α, the greater is the deviation from Bayesian updating, the larger is the temptation to panic and the more likely is it that the agent will yield to the urge to leave stocks entirely.

Confirmatory Bias: In the model with Positive Prior-Bias ($\lambda \geq 0$), one can interpret (3.16) as describing an agent who overlooks the evidence represented by s_1 with probability $\gamma = \frac{\alpha\lambda}{1+\alpha}$. This interpretation is valid even if γ varies with the signal. When $\gamma(s_1)$ varies suitably with s_1, one obtains a model where the probability of overlooking a particular piece of evidence (and retaining prior beliefs) is smaller when the evidence supports prior beliefs. Such a bias towards supportive evidence is reminiscent of the well-documented *confirmatory bias*; see

Rabin and Schragg [19] for references to the relevant psychology literature and for an alternative model of the bias.[17]

To illustrate, suppose that

$$S_1 = \{a, b\},\ S_2 = \{A, B\},\ \text{ and }\ p\left(a \mid A\right) = p\left(b \mid B\right) > \tfrac{1}{2}. \tag{3.17}$$

Then B is more likely under prior beliefs ($p_2\left(B\right) > \frac{1}{2}$) if and only if $p_1\left(b\right) > \frac{1}{2}$. Thus the desired bias is captured by the specification

$$\gamma\left(s_1\right) = \gamma^*\left(\frac{p_1\left(s_1\right)}{max_{s_1' \in S_1} p_1\left(s_1'\right)}\right), \quad s_1 = a, b,$$

with $\gamma^* : [0,1] \longrightarrow [0,1]$ decreasing. If she believes initially that B is more likely than A, then the conflicting signal a will be overlooked with high probability.

Representativeness: Once again, adopt (3.17). The likelihood information given there indicates that a is representative of A and b is representative of B. According to the representative heuristic (Tversky and Kahneman [21], for example), people often weight such representativeness too heavily when judging conditional probabilities of A given a and B given b. To capture the resulting updating bias, take

$$q\left(A \mid a\right) = q\left(B \mid b\right) = 1.$$

Then the conditional of the compromise prior given by (3.7) satisfies

$$p^*\left(A \mid a\right) > p\left(A \mid a\right) \text{ and } p^*\left(B \mid b\right) > p\left(B \mid b\right).$$

Sample-Bias: Think of repeated trials of an experiment and take $S_1 = S_2 = S$. Denote by $\delta_s\left(\cdot\right)$ the measure assigning probability 1 to s and let

$$q\left(\cdot \mid s\right) = \left(1 - \lambda\right)\ p\left(\cdot \mid s\right) + \lambda\, \delta_s\left(\cdot\right),$$

where $\lambda \leq 1$ is a constant. When $\lambda > 0$, the Bayesian update of p is adjusted in the direction of the "empirical frequency" measure $\delta_s\left(\cdot\right)$, implying a bias akin

[17] Rabin and Schragg's principal model assumes that with positive probability the agent misinterprets an unfavorable signal as supporting prior beliefs. Here an unfavorable signal is recognized as such but is underweighted. They suggest the latter as a reasonable alternative.

to the *hot-hand fallacy* - the tendency to overpredict the continuation of recent observations. If $\lambda < 0$, then[18]

$$q(\cdot \mid s) = p(\cdot \mid s) - \lambda \ (p(\cdot \mid s) - \delta_s(\cdot)),$$

and the adjustment is proportional to $(p(\cdot \mid s) - \delta_s(\cdot))$, as though expecting the next realization to compensate for the discrepancy between $p(\cdot \mid s)$ and the most recent observation. This is a form of negative correlation with past realizations akin to the *gambler's fallacy.*

4. CONCLUDING REMARKS

The connection of our model to updating relies on the interpretation of the functional form (3.5) as describing interim choice at time 1. This interpretation is suggested by our formal model, but Theorem 3.1 deals only with the ex ante choice between contingent menus and not with the interim choice of acts from menus. A similar issue arises in the GP model and their solution, using suitably extended preferences, can be adapted here. Note that foundations provided in this way are subject to the difficulty pointed out in [10, p. 1415], namely the lack of a revealed preference basis for extended preferences.

One might expect a given individual to update differently in different situations. If by 'situation' one means 'state space', then the present model is consistent with such variation because it is restricted to a given state space. However, it has nothing to say about how behavior is connected across state spaces. Alternatively, one might expect that even given the state space, an individual may exhibit different updating biases depending on the choice problem. This calls for a generalization that would permit updating behavior to depend on the menu available at the interim stage.

Conclude with one application of the model. Our non-Bayesian agent violates the law of iterated expectations, because she uses the compromise measure p^* from (3.6) to guide choice at $t = 1$ but she uses p for choice at $t = 0$. Thus there exist acts $f : S_2 \longrightarrow X$ such that $\{f\} \succ \{-f\}$ at time 0 and yet such that at each s_1, the agent would strictly prefer to choose $-f$ out of $\{f, -f\}$. This amounts to a violation of the "sure-thing principle for action rules", a property that has been identified as central to no-trade theorems, (see Geanakoplos [8, Section 6], for

[18] To ensure that $q(\cdot \mid s)$ is a probability measure (hence non-negative), assume that $p(s \mid s) \geq \frac{-\lambda}{1-\lambda}$ for all s.

example). It is not surprising, therefore, that two such agents may agree to take opposite sides of a bet at every state s_1 even if they have common commitment priors and all the preceding is common knowledge. What may be not so obvious, however, is that common knowledge agreement to bet can arise even though agents have stable (unchanging) preferences, albeit over contingent menus rather than on the usual domain of acts. Future research will explore more deeply the message of this example - that trade may result not only from heterogeneity in prior beliefs or in information, but also from heterogeneity in the way that agents update in response to information.

A. APPENDIX: Proof of Representation Result

Necessity: Set-Betweenness is satisfied because each $\mathcal{U}(\cdot; s_1)$ defined in (3.3) has the GP form. The pair (s_1, s_2) is null iff it is strongly null iff $p(s_1, s_2) = 0$; use $q << p$ to prove sufficiency of the latter. To verify RSR, note that for any menu of lotteries L,

$$\mathcal{U}(Ls_2M; s_1) =$$

$$\max_{f \in Ls_2M} \left\{ \int_{S_2} u(f)\, dp(\cdot \mid s_1) + \alpha(s_1) \int_{S_2} u(f)\, dq(\cdot \mid s_1) \right\}$$

$$- \max_{f' \in Ls_2M} \alpha(s_1) \int_{S_2} u(f')\, dq(\cdot \mid s_1)$$

$$= \max_{f \in M} \left\{ \int_{S_2 \setminus \{s_2\}} u(f)\, dp(\cdot \mid s_1) + \alpha(s_1) \int_{S_2 \setminus \{s_2\}} u(f)\, dq(\cdot \mid s_1) \right\}$$

$$- \max_{f' \in M} \alpha(s_1) \int_{S_2 \setminus \{s_2\}} u(f')\, dq(\cdot \mid s_1)$$

$$+ \max_{\ell \in L} \{u(\ell)\, p(s_2 \mid s_1) + \alpha(s_1)\, u(\ell)\, q(s_2 \mid s_1)\} - \max_{\ell' \in L} \alpha(s_1)\, u(\ell')\, q(s_2 \mid s_1).$$

Because only the last two terms depend on L, for the purpose of describing the induced ranking of menus of lotteries, we can write $\mathcal{U}(Ls_2M; s_1) =$

$$\max_{\ell \in L} u(\ell)\, \{p(s_2 \mid s_1) + \alpha(s_1)\, q(s_2 \mid s_1)\} - \max_{\ell' \in L} \alpha(s_1)\, u(\ell')\, q(s_2 \mid s_1)$$

$$= p(s_2 \mid s_1) \max_{\ell \in L} u(\ell),$$

and this implies RSR. Other axioms are readily verified.

Sufficiency: The argument is straightforward. By adapting arguments from Anscombe-Aumann and Gul-Pesendorfer, the first seven axioms imply the representation (3.1)-(3.2). Though GP deal with menus of lotteries while we have menus of Anscombe-Aumann acts over S_2, their representation result [10, Theorem 3] can be translated into our setting and applied. Next, RSR and Absolute Continuity deliver (3.3) and Prior-Bias delivers (3.8).

Because $\succeq$ satisfies Order, Continuity and Independence on $\mathcal{C}$, there exists a representation for $\succeq$ of the form

$$\mathcal{U}(F) = \Sigma_{s_1 \in S_1} \mathcal{U}(F(s_1); s_1), \tag{A.1}$$

where each $\mathcal{U}(\cdot; s_1)$ is continuous and mixture linear on $\mathcal{M}(S_2)$. Intuition for this claim is provided by the similarity with the Anscombe-Aumann theorem. The latter deals with acts mapping a state space into $\Delta(X)$, while here each F maps states into $\mathcal{M}(S_2)$, which shares with $\Delta(X)$ the existence of a mixing operation. However, $\mathcal{M}(S_2)$ is not a mixture space and thus the analogy is not perfect.[19] Here are the missing details: Adapt [10, Lemma A.1] and define $\succeq^*$ on $\Delta(\mathcal{C})$, the set of finite support lotteries on $\mathcal{C}$, by

$$\pi' \succeq^* \pi \iff \Sigma_{F \in \mathcal{C}} \pi'(F) F \succeq \Sigma_{F \in \mathcal{C}} \pi(F) F, \quad \text{for all } \pi', \pi \in \Delta(\mathcal{C}).$$

Then the vNM theorem delivers $\mathcal{U} : \mathcal{C} \longrightarrow \mathbf{R}^1$ such that $\pi \longmapsto \Sigma_{F \in \mathcal{C}} \pi(F) \mathcal{U}(F)$ represents $\succeq^*$. It follows that $\mathcal{U}$ represents $\succeq$ on $\mathcal{C}$ and is linear and continuous there. Now apply Kreps [13, Propn. 7.4], specifically the proof of his equation (7.10), to derive the additive representation (A.1).[20]

It follows that for each s_1, $\succeq$ induces a conditional order $\succeq_{s_1}$ on $\mathcal{M}(S_2)$ and that the latter is represented by $\mathcal{U}(\cdot; s_1)$. By Set-Betweenness, $\succeq_{s_1}$ satisfies the GP axioms suitably translated to our setting; that is, GP deal with menus of lotteries, while we have menus of Anscombe-Aumann acts over S_2. With this translation, their proof [10, Theorem 3] is valid for our setting and delivers:

$$\mathcal{U}(M; s_1) = \max_{f \in M} \left\{ U(f; s_1) + V(f; s_1) - \max_{f' \in M} V(f'; s_1) \right\}, \tag{A.2}$$

where $U(\cdot; s_1)$ and $V(\cdot; s_1)$ are mixture linear (and continuous). Actually, the preceding equality is valid only up to ordinal equivalence, but both sides are

[19] It violates the property $\lambda(\lambda' M + (1 - \lambda') N) + (1 - \lambda) N = \lambda\lambda' M + (1 - \lambda\lambda') N$.

[20] Kreps deals with acts f for which each lottery $f(s)$ has finite support, but the extension to all f is valid under Continuity.

mixture linear and so they must be cardinally equivalent. Thus absolute equality may be assumed wlog.

Both $U(\cdot; s_1)$ and $V(\cdot; s_1)$, utility functions defined on the domain $(\Delta(X))^{S_2}$ of Anscombe-Aumann acts, satisfy the basic mixture space axioms there. Thus, again by [13, Propn. 7.4] and Continuity, we can write

$$U(f; s_1) = \Sigma_{s_2 \in S_2} u(f(s_2); s_1, s_2), \ V(f; s_1) = \Sigma_{s_2 \in S_2} v(f(s_2); s_1, s_2), \quad \text{(A.3)}$$

for all $f : S_2 \longrightarrow \Delta(X)$. (Below we often abbreviate (s_1, s_2) by s.) Each $u(\cdot; s)$ and $v(\cdot; s)$ is mixture linear and continuous.

Next apply State Independence. Take $F \in \mathcal{A}$, that is, let $F(s_1)$ be a singleton for every s_1. Recall that $\mathcal{A}$ is isomorphic to $(\Delta(X))^{S_1 \times S_2}$, the set of Anscombe-Aumann acts over $S_1 \times S_2$. The preceding three displayed equations imply that $\succeq$ restricted to $\mathcal{A}$ is represented by $\widehat{U}$, where

$$\widehat{U}\left(\widehat{f}\right) = \Sigma_{s_1, s_2} u\left(\widehat{f}(s_1, s_2); s_1, s_2\right), \text{ for all } \widehat{f} \in (\Delta(X))^{S_1 \times S_2}.$$

By State Independence and Nondegeneracy, the order represented by $\widehat{U}(\cdot)$ satisfies all the Anscombe-Aumann axioms. Thus (by [13, Theorem 7.17], for example) $\widehat{U}$ has the SEU form

$$\widehat{U}\left(\widehat{f}\right) = \Sigma_{s \in S_1 \times S_2} p(s) \, u\left(\widehat{f}(s)\right),$$

for suitable u and probability measure p on $S_1 \times S_2$. (Equality is modulo ordinal equivalence, but the latter qualification can be dropped because $\widehat{U}(\cdot)$ is mixture linear, forcing the ordinal transformation to be cardinal.) Wlog therefore,

$$u(\cdot; s_1, s_2) = p(s_1, s_2) \, u(\cdot)$$

and we can refine (A.3) and write

$$U(f; s_1) = \Sigma_{s_2 \in S_2} p(s_1, s_2) \, u(f(s_2)). \quad \text{(A.4)}$$

The next step is to show that on X,

$$v(\cdot; s) = a_s p(s) u(\cdot) + b_s, \text{ where } a_s \geq 0. \quad \text{(A.5)}$$

To do so, note that for any $s = (s_1, s_2)$ and menus $M \in \mathcal{M}(S_2)$ and $L \subset \Delta(X)$,

$$\mathcal{U}(L s_2 M; s_1) = \max_{f \in M, \ell \in L} \{U(\ell s_2 f; s_1) + V(\ell s_2 f; s_1)\}$$

$$- \max_{f' \in M, \ell' \in L} V\left(\ell' s_2 f'; s_1\right) \text{ , and hence}$$

$$\mathcal{U}\left(L s_2 M; s_1\right) = \max_{\ell \in L} \left\{ p\left(s\right)\, u\left(\ell\right) + v\left(\ell; s\right) - \max_{\ell' \in L} v\left(\ell'; s\right) \right\} + \Phi\left(M, s\right), \quad \text{(A.6)}$$

where the final term is independent of L and thus can be ignored in what follows. By State Independence, the invariance in RSR applies to all non-null (s_1, s_2), and hence also to all (s_1, s_2). Thus, by RSR, the ranking of menus of lotteries represented by $L \longmapsto \mathcal{U}\left(L s_2 M; s_1\right)$ is strategically rational.

Case 1: Suppose $p\left(s\right) > 0$. Then $p\left(s\right) u\left(\cdot\right)$ is nonconstant. Hence (A.5) follows as in [10, p. 1414].

Case 2: Suppose that $p\left(s\right) = 0$. Then (A.6) and Absolute Continuity imply that the utility of any F is independent of what it assigns to the state s. Thus *any* specification for $v(\cdot; s)$ is consistent with a representation for $\succeq$. In particular, we can take $v\left(\cdot; s\right) = 0$ wlog and (A.5) is valid with $a_s = 0$.

From (A.3)-(A.5), deduce that

$$\mathcal{U}\left(M; s_1\right) = \max_{f \in M} \left\{\Sigma_{s_2}\, p\left(s_1, s_2\right) u\left(f\left(s_2\right)\right) + \Sigma_{s_2} a_{s_1, s_2} p\left(s_1, s_2\right) u\left(f\left(s_2\right)\right)\right\}$$

$$- \max_{f' \in M} \left\{\Sigma_{s_2} a_{s_1, s_2} p\left(s_1, s_2\right) u\left(f'\left(s_2\right)\right)\right\}.$$

Let

$$\alpha\left(s_1\right) = \frac{\Sigma_{s_2} a_{s_1, s_2} p\left(s_1, s_2\right)}{p_1\left(s_1\right)} \geq 0.$$

(p_1 is the S_1-marginal of p; it is everywhere positive by (A.3), (A.5) and S_1-Non-nullity.) Define the measure q so that its S_1-marginal equals p_1, and

$$q\left(s_2 \mid s_1\right) = \begin{cases} \frac{a_{s_1, s_2} p(s_2 \mid s_1)}{\alpha(s_1)} & \text{if } \alpha\left(s_1\right) > 0 \\ p\left(s_2 \mid s_1\right) & \text{otherwise.} \end{cases}$$

Then

$$\mathcal{U}\left(M; s_1\right) = p_1\left(s_1\right) \max_{f \in M} \left\{\Sigma_{s_2}\, p\left(s_2 \mid s_1\right) u\left(f\left(s_2\right)\right) + \alpha\left(s_1\right)\, \Sigma_{s_2} q\left(s_2 \mid s_1\right) u\left(f\left(s_2\right)\right)\right\}$$

$$- p_1\left(s_1\right) \max_{f' \in M} \left\{\alpha\left(s_1\right)\, \Sigma_{s_2} q\left(s_2 \mid s_1\right) u\left(f'\left(s_2\right)\right)\right\}.$$

With (A.1), this yields the desired representation (3.1)-(3.2).

Finally, assume also Prior-Bias. Given the representation just established, the axiom can be translated into the statement:

$$\text{if } \int (u(f) - u(g))\, dp(\cdot \mid s_1) > 0 \text{ and } \int (u(f) - u(g))\, dp_2(\cdot) = 0,$$

$$\text{then } \alpha(s_1) \int (u(f) - u(g))\, dq(\cdot \mid s_1) \geq 0.$$

The desired result now follows by a Theorem of the Alternative [16, **?**].

Similarly for Positive and Negative Prior-Bias. ■

Proof of Corollary 3.3: If (u, p, q, α) represents $\succeq$ and if (u', p', q', α') is related as stated, then clearly it also represents $\succeq$. For the converse, suppose that both tuples represent $\succeq$. Then the subjective expected utility functions defined by (u, p) and (u', p') both represent preference on the subset $\mathcal{A} \subset \mathcal{C}$ of Anscombe-Aumann acts over $S_1 \times S_2$. By the well-known uniqueness properties of the Anscombe-Aumann theorem,

$$u' = au + b \text{ and } p' = p. \tag{A.7}$$

By Lemma 3.2, (u', p', q', α') violates (3.11) iff (u, p, q, α) does, in which case (3.12) is valid. Suppose that (u, p, q, α) satisfies (3.11). The latter implies that $\mathcal{U}(\cdot; s_1)$ is regular in the sense of GP (p. 1414). (Here and below we refer to the translation of GP to our setup, whereby their menus of lotteries are replaced by menus of Anscombe-Aumann acts over S_2.) Thus their Theorem 4 implies that

$$V'(\cdot; s_1) = A_{s_1} V(\cdot; s_1) + B_{V, s_1} \quad \text{and} \tag{A.8}$$

$$U'(\cdot; s_1) = A_{s_1} U(\cdot; s_1) + B_{U, s_1}, \tag{A.9}$$

where

$$V(f; s_1) = \alpha(s_1) \int_{S_2} u(f)\ dq(\cdot \mid s_1), \qquad f \in (\Delta(X))^{S_2}, \tag{A.10}$$

and V', U', V and U are defined similarly. Deduce from (A.7) and (A.9) that $A_{s_1} = a$. Equation (A.8) implies, again by uniqueness properties of the Anscombe-Aumann model, that $q'(\cdot \mid s_1) = q(\cdot \mid s_1)$. Substitution from (A.10) implies further that

$$\alpha'(s_1) \int_{S_2} (au(f) + b)\ dq(\cdot \mid s_1) = a\, \alpha(s_1) \int_{S_2} u(f)\ dq(\cdot \mid s_1) + B_{V, s_1},$$

for all $f \in (\Delta(X))^{S_2}$, which implies that $\alpha'(s_1) = \alpha(s_1)$. ■

B. APPENDIX: Changing Risk Aversion

This appendix elaborates on the model of changing risk preference mentioned following Theorem 3.1. More precisely, consider the utility function on $\mathcal{C}$ given by

$$\mathcal{U}(F) = \int_{S_1} \mathcal{U}(F(s_1); s_1)\, dp_1, \quad F \in \mathcal{C}, \tag{B.1}$$

and $\mathcal{U}(F(s_1); s_1) =$

$$\left\{ \max_{f \in F(s_1)} \int_{S_2} [u(f) + v(f)]\, dp(\cdot \mid s_1) \right\} - \max_{f' \in F(s_1)} \int_{S_2} v(f')\, dp(\cdot \mid s_1), \tag{B.2}$$

where:

*Reg1** $u, v : \Delta(X) \longrightarrow \mathbb{R}^1$ are mixture linear and continuous.

*Reg2** p_1 has full support on S_1; each $p(\cdot \mid s_1)$ is a probability measure on S_2.

*Reg3** There exist lotteries ℓ' and ℓ such that $v(\ell') < v(\ell)$ and $u(\ell') + v(\ell') > u(\ell) + v(\ell)$.

*Reg4** u cannot be expressed as the linear transformation $u = av + b$, where $a < 0$.

Note that these conditions rule out the standard Bayesian model (say if v is constant or a positive affine transformation of u) and that the above model is disjoint from our central model. Turn to the axioms that underlie them.

Adopt all previous axioms with the exception of RSR. To state its replacement, define the induced order $\succeq_{(s_1,s_2)}$ by: for all menus L' and L of lotteries,

$$L' \succeq_{(s_1,s_2)} L \text{ iff } (F_{-s_1}, L's_2M) \succeq (F_{-s_1}, Ls_2M),$$

for some F and M. Because GP also deal with preferences on menus of lotteries, we can adapt their terminology. Say that $\succeq_{(s_1,s_2)}$ has self-control if $L' \succ_{(s_1,s_2)} L' \cup L \succ_{(s_1,s_2)} L$ for some menus of lotteries. Refer to $\succeq_{(s_1,s_2)}$ as having *regular self-control* if it has self-control and there exists L such that: (i) $L \succeq_{(s_1,s_2)} L'$ for all subsets L' of L, and (ii) $L \succ_{(s_1,s_2)} \{\ell\}$ for some ℓ in L. GP (p. 1414) refer to (i)-(ii) as the absence of maximal preference for commitment.[21]

[21] They also define regularity. The connection is that, given that $\succeq_{(s_1,s_2)}$ has self-control, then it has regular self-control iff it is regular.

Axiom 13 (Regular Self-Control). *There exists a state* (s_1, s_2) *such that* $\succeq_{(s_1,s_2)}$ *has regular self-control.*

A state satisfying the conditions of the axiom is necessarily non-null. Therefore, given State Independence, "there exists (s_1, s_2)" can be replaced by "for every non-null (s_1, s_2)."

To interpret the axiom, compare it with RSR. Roughly, the latter says that where the choice is between contingent menus that differ only in one state, then there is no temptation, while the new axiom says the opposite - that temptation does sometimes occur even for such choices. Thus while RSR implies that temptation arises only where beliefs are relevant for choice, Regular Self-Control implies that it arises even where only risk attitude is relevant for choice. This explains why substituting Regular Self-Control for RSR leads to the following alternative to Theorem 3.1.

Theorem B.1. $\succeq$ *satisfies Order, Continuity, Independence, Nondegeneracy,* S_1*-Non-nullity, Set-Betweenness, State Independence, Regular Self-Control and Absolute Continuity if and only if it admits a representation of the form (B.1)-(B.2), including the regularity conditions Reg1*-Reg4*.*

Though Theorems 3.1 and B.1 are disjoint, they do not exhaust the class of preferences satisfying all their common axioms. As an example, take (B.1) and $\mathcal{U}(F(s_1); s_1) =$

$$\left\{\max_{f\in F(s_1)} \int_{S_2} u(f)\ dp(\cdot \mid s_1) - \alpha \int_{S_2} u(f)\ dq(\cdot \mid s_1)\right\} - \max_{f'\in F(s_1)} (-\alpha) \int_{S_2} u(f)\ dq(\cdot \mid s_1)$$

$$= \left\{\max_{f\in F(s_1)} \int_{S_2} u(f)\ dp(\cdot \mid s_1) - \alpha \int_{S_2} u(f)\ dq(\cdot \mid s_1)\right\} + \min_{f'\in F(s_1)} \alpha \int_{S_2} u(f)\ dq(\cdot \mid s_1),$$

where $q(\cdot \mid s_1) << p(\cdot \mid s_1)$ and $1 < \alpha$. Though all other axioms are satisfied, RSR and Regular Self-Control are violated and thus the functional form does not fit into either model. The interpretation is that the underlying "change in preference" cannot be attributed to a change in only one of taste or beliefs.[22]

Proof. Necessity: Regular Self-Control is implied by Reg3*, Reg4* and the observations in GP (p. 1414).

[22] The obvious funtional form based on arbitrary pairs (u, p) and (v, q) would also support such a statement, but it would violate State Independence.

Sufficiency: As in the proof of Theorem 3.1, we have (A.1)-(A.4). Our objective is to prove that wlog

$$v(\cdot; s_1, s_2) = p(s_1, s_2)\, \widehat{v}(\cdot) \quad \text{for all } (s_1, s_2). \tag{B.3}$$

Conditions Reg1*-Reg4* would then follow; for example, Reg4* would follow from Regular Self-Control along the lines of the observations in GP (p. 1414).

Let (s_1^*, s_2^*) satisfy the conditions in Regular Self-Control. Then (s_1^*, s_2^*) is non-null. By State Independence, $\succeq_{(s_1,s_2)} = \succeq_{(s_1^*,s_2^*)}$ for any other non-null (s_1, s_2). Thus $\succeq_{(s_1,s_2)}$ is represented by both

$$W(L; s_1, s_2) = max_{\ell \in L}\{p(s_1, s_2)\, u(\ell) + v(\ell; s_1, s_2)\} - max_{\ell \in L} v(\ell; s_1, s_2),$$

and by the corresponding function $W(\cdot; s_1^*, s_2^*)$. Moreover, $\succeq_{(s_1,s_2)}$ satisfies the conditions in [10, Theorem 4]. Thus

$$p(s_1, s_2)\, u(\cdot) = ap(s_1^*, s_2^*)\, u(\cdot) + b_u,\; v(\cdot; s_1, s_2) = av(\cdot; s_1^*, s_2^*) + b_v, \tag{B.4}$$

for some common $a > 0$ (that depends on the two states). From the first equation,

$$(p(s_1, s_2) - ap(s_1^*, s_2^*))\; u(\cdot) \text{ is constant,}$$

which implies that

$$p(s_1, s_2) - ap(s_1^*, s_2^*) = 0 = b_u.$$

Therefore,

$$v(\cdot; s_1, s_2) = \tfrac{p(s_1,s_2)}{p\left(s_1^*,s_2^*\right)}\, v(\cdot; s_1^*, s_2^*), \tag{B.5}$$

where wlog b_v has been set equal to zero. It is easy to see from above and Absolute Continuity, that $p(s_1, s_2) > 0$ if (s_1, s_2) is non-null. Thus $\frac{v(\cdot;s_1,s_2)}{p(s_1,s_2)}$ is invariant across non-null states. Denote the common function by $\widehat{v}(\cdot)$, which yields (B.3) for all non-null states s.

Finally, if (s_1, s_2) is null, then the utility of any F is independent of what it assigns to (s_1, s_2). Thus *any* specification for $v(\cdot; s_1, s_2)$ is consistent with a representation for $\succeq$. In particular, we can take $v(\cdot; s_1, s_2) = 0$ consistent with (B.3). ■

References

[1] C.D. Aliprantis and K.C. Border, *Infinite Dimensional Analysis*, Springer, 1994.

[2] S. Barbera and B. Grodal, Preference for flexibility and the opportunities of choice, Institute of Economics 02-22, U. Copenhagen, 2002.

[3] M. Brandt, Q. Zeng and L. Zhang, Equilibrium stock return dynamics under alternative rules of learning about hidden states, *J. Econ. Dyn. and Control*, forthcoming.

[4] A. Brav and J.B. Heaton, Competing theories of financial anomalies, *Rev. Finan. Stud.* 15 (2002), 575-606.

[5] C. Camerer, Individual decision-making, in *Handbook of Experimental Economics*, J. Kagel and A. Roth eds., Princeton U. Press, 1995.

[6] E. Dekel, B. Lipman and A. Rustichini, A unique subjective state space for unforeseen contingencies, *Econometrica* 69 (2001), 891-934.

[7] L.G. Epstein and M. Schneider, Recursive multiple-priors, *J. Econ. Theory* 113 (2003), 1-31.

[8] J. Geanakoplos, Common knowledge, in *Handbook of Game Theory* vol. 2, Aumann and Hart eds., Elsevier, 1994.

[9] P. Ghirardato, Revisiting Savage in a conditional world, *Econ. Theory* 20 (2002), 83-92.

[10] F. Gul and W. Pesendorfer, Temptation and self-control, *Econometrica* 69 (2001), 1403-35.

[11] F. Gul and W. Pesendorfer, Self-control and the theory of consumption, *Econometrica* 72 (2004), 119-158..

[12] D. Kreps, A representation theorem for 'preference for flexibility', *Econometrica* 47 (1979), 565-577.

[13] D. Kreps, *Notes on the Theory of Choice*, Westview, 1988.

[14] 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.

[15] D. Laibson, Golden eggs and hyperbolic discounting, *Quart. J. Econ.* 12 (1997), 443-477.

[16] O.L. Mangasarian, *Nonlinear Programming*, McGraw-Hill, New York: 1969.

[17] K. Nehring, Preference for flexibility in a Savage framework, *Econometrica* 67 (1999), 101-120.

[18] M. Rabin, Psychology and economics, *J. Ec. Lit.* 36 (1998), 11-46.

[19] M. Rabin and J. Schragg, First impressions matter: a model of confirmatory bias, Quart. J. Econ. (1999), 37-82.

[20] R.H. Strotz, Myopia and inconsistency in dynamic utility maximization, *Rev. Ec. Studies* 23 (1955), 165-180.

[21] A. Tversky and D. Kahneman, Judgement under uncertainty: heuristics and biases, *Science* 185 (1974), 1124-31.