

Inertial Behavior and Generalized Partition

David
Dillenberger

University of
Pennsylvania

Philipp
Sadowski

Duke University

May 1, 2016

ERID Working Paper Number 216

This paper can be downloaded without charge from the Social
Science Research Network Electronic Paper Collection:

<http://ssrn.com/abstract=2787996>

**Economic Research Initiatives at Duke
WORKING PAPERS SERIES**



Inertial Behavior and Generalized Partition*

David Dillenberger[†] Philipp Sadowski[‡]

May 2016

Abstract

Behavior is inertial if it does not react to the apparent arrival of relevant information. In a context where information is subjective, we formulate an axiom that captures inertial behavior, and provide a representation that explains such behavior as that of a rational decision maker who perceives a particular type of information structure, we call a generalized partition. We characterize the learning processes that can be described by a generalized partition. We then investigate behavioral protocols that may lead individuals to perceive a generalized partition (and thus to display inertial behavior) even when facing a different type of information structure: A cognitive bias referred to as cognitive inertia and a bound on rationality, which we term shortsightedness.

Key words: Inertial behavior, subjective learning, generalized partition, uniform cover, cognitive inertia, shortsighted rationality

1. Introduction

Individuals often seem reluctant to change their behavior in response to new information, unless that information is conclusive. Examples include managers who do not adapt their strategy to changes in the business environment (Hodgkinson 1997 and reference therein), a bias in favor of renewing the status quo health plan when reviewing other options after a change in health status, or a “one advice suits all” approach among consultants, such as medical doctors or financial advisors, dealing with a heterogeneous population (clients with different symptoms or different levels of risk attitude). We refer to this apparent bias as inertial behavior.

We contend that in most applied situations, the analyst may suspect inertial behavior even without being aware of the exact information available to each agent; behavior does

*Some of the results in this paper previously appeared in Dillenberger, D., and P. Sadowski (2012). We thank David Ahn, Brendan Daley, Jason Dana, Haluk Ergin, Itzhak Gilboa, Faruk Gul, Peter Landry, Wolfgang Pesendorfer, Todd Sarver, Andrei Savochkin, Marciano Siniscalchi, Roe Teper, and especially Kfir Eliaz for their comments and suggestions.

[†]Department of Economics, University of Pennsylvania. E-mail: ddill@sas.upenn.edu

[‡]Department of Economics, Duke University. E-mail: p.sadowski@duke.edu

not change although it is likely that some inconclusive information became available. But if information is subjective, then it is possible that inertial behavior is not actually biased (compared to the rational Bayesian benchmark), but instead is the optimal course of action given the information structure perceived by the agent.

Our goal in this paper is to analyze sources of inertial behavior under the assumption that learning about the relevant state of the world is subjective. We start by eliciting the information structure a decision maker (henceforth DM) perceives from his observed choice behavior. We then formulate an axiom that captures inertial behavior, and provide a representation that explains such behavior as that of a rational DM who perceives a particular type of information structure, called a *generalized partition*. While consistent with Bayesian updating, this representation does not rule out the possibility that the DM nevertheless ignores some relevant information and behaves sub-optimally. To better understand where inertial behavior could originate from, we proceed to assume that there is a true underlying information structure, and investigate different channels that can lead the DM to perceive this information structure as a generalized partition (and thus to display inertial behavior).

In order to elicit the subjective information structure from observed behavior, we follow the approach in Dillenberger, Lleras, Sadowski, and Takeoka (2014, henceforth DLST) and take as primitive a preference relation over sets (or menus) of acts, which are defined over a given space of payoff-relevant states, S .^{1,2} This allows us to identify a suitable notion of anticipated preferences over acts contingent on learning an event. Our notion of inertial behavior then connects the ex ante comparison of acts (i.e., singleton menus) to those anticipated ex post preferences. Loosely speaking, the *Inertial Behavior* axiom requires that if $I \subseteq S$ is an event the DM can foresee learning, then he prefers to commit to act $fI0$, which agrees with f on I and pays zero otherwise, over committing to $gI0$, only if he also anticipates to prefer f to g contingent on actually learning I .

Theorem 1 derives a *generalized-partition representation*, which can be interpreted as follows: the DM behaves as if he (i) holds prior beliefs over the state space; and (ii) perceives a particular type of stochastic rule, referred to as a *generalized partition*, that determines which event he will learn contingent on the true state. Upon learning an event, the DM calculates posterior beliefs using Bayes' law, which leads him to exclude all states that are not in that event, keeping the relative likelihood of the remaining states fixed (a property

¹The interpretation is that the DM initially chooses among menus and subsequently chooses an act from the menu. If the ultimate choice of an act takes place in the future, then the DM may expect information to arrive prior to this choice. Analyzing preferences over future choice situations allows one to identify the anticipated future choice behavior without observing ex post preferences, and to interpret it in terms of the information the DM expects to receive.

²The set S may represent the state space on which acts are naturally defined in a particular application. The same primitive is also analyzed in de Oliveira et al. (2016).

we term *non-shifting-weights*). He then chooses from the menu the act that maximizes the corresponding expected utility. The prior and the generalized partition are uniquely identified.

Turning to discuss the different channels that might lead the DM to perceive a generalized partition, the most immediate one is that the true underlying information structure is itself a generalized partition. We show that the class of generalized partitions can accommodate a variety of information structures that correspond to intuitive learning processes and to plausible behavior. As the name suggests, generalized partition extends the notion of set partition, according to which the DM learns which cell of a partition contains the true state. In the case of a set partition, signals are deterministic; that is, for each state there is only one possible event that contains it. Another example of a generalized partition is a random partition, where one of multiple partitions is randomly drawn and then an event in it is reported. A situation that may give rise to a random partition is an experiment with uncertainty about its precision. A sequential elimination of candidates, say during a recruiting process, where k out of n candidates are to be eliminated in the first stage (so that the resulting collection of events the DM might learn is the set of all $(n - k)$ -tuples), can also be modeled by an information structure we can accommodate. Theorem 2 characterizes all types of learning that can be described by a generalized partition.

Alternatively, as we discuss in Section 4.1, the DM could indeed be biased and display *cognitive inertia*, in the sense that he ignores all information that is not conclusive. That is, even when receiving signals that speak to the relative likelihood of different states, he believes that the content of any signal is merely the collection of states that have been completely ruled out by it. Our model is consistent with this explanation, but establishes that it, too, is restrictive: the true information structure must be such that ignoring all inconclusive information leads to a generalized partition.

Lastly, we propose a particular bound on rationality, which leads the DM to behave as if he faces a generalized partition for any true underlying information structure. Our bound is inspired by the *construal level theory* in social psychology, according to which the closer in time (or physical distance) a situation is from the individual, the more easily and concretely it will be thought of. To wit, decisions about the distant future, for example about careers after graduation or life after retirement, are challenging not only because relevant information is noisy far in advance, but simply because individuals cannot think concretely about contingencies that may occur in the distant future.³

³This is conceptually different from the idea of present bias. According to the later, individuals may trade off outcomes in consecutive periods differently, depending on whether those periods lie in the immediate or distant future. In contrast, the boundedly rational individual we consider finds it hard to make any comparison between the two outcomes if they are relevant only in the distant future.

To capture this intuition, we propose a bound on the horizon up to which the DM can independently form beliefs about his uncertain environment, and call such a DM *shortsighted*. Suppose nature determines the state in two stages. In the first stage some states are ruled out, and in the second stage all but one of the remaining ones are removed. The DM is able to directly form beliefs only over the first stage; that is, he can not independently form beliefs over the further resolution of uncertainty in the second stage. This constraint has two consequences. First, the DM does not initially hold a prior over final states, but only over the events he might learn in the first period. Second, he is oblivious to any correlation in the resolution of uncertainty across periods.⁴

It seems plausible that, given this bound, a rational DM will try to construct a prior on S , as well as an information structure that features no correlation across periods (i.e., a generalized partition),⁵ so that jointly they match his beliefs for the first period. Theorem 3 establishes that this construction is always possible, while Theorem 4 provides a sufficient condition for the constructed prior and generalized partition to be unique. In other words, a shortsighted, but otherwise rational, DM can always be described as if he faces a generalized partition.

The remainder of the paper is organized as follows. Section 2 formalizes our notion of inertial behavior and provides our representation result. Section 3 characterizes the learning processes that generate a generalized partition. Section 4 provides the two protocols that lead to learning via a generalized partition even if the true underlying information structure does not belong to this class. Related literature is discussed in Section 5. Most proofs are relegated to the appendices.

2. Inertial Behavior and Generalized Partition

Our analysis is broken into three parts. First, we discuss which information structures give rise to inertial behavior under Bayesian updating. Second, we formulate an axiom that captures inertial behavior in terms of preferences over menus of acts when the information structure, as well as the ex post preferences, are unobservable. Finally, our main theorem establishes that, when information is subjective, inertial behavior of a forward looking DM can always be explained as that of a Bayesian who faces an appropriate information structure, namely a generalized partition.

⁴We identify signals with events here. This is without loss, because two different signals that have the same support will be equivalent for a DM who does not understand their correlation with the states in that support.

⁵Even if the DM were aware that there *could* be correlation across periods, the principle of insufficient reason might lead him to focus on generalized partitions.

2.1. Inertial Behavior and the Non-Shifting Weights Property

Let $S = \{s_1, \dots, s_k\}$ be a finite state space. An act is a mapping $f : S \rightarrow [0, 1]$. We interpret payoffs in $[0, 1]$ to be in utils; that is, we look directly at utility acts (assuming that the cardinal utility function over outcomes is known and payoffs are stated in its units) instead of the corresponding Anscombe-Aumann (1963) acts. For any number $c \in [0, 1]$, we simply denote by c the constant act that yields c in every state. Let \mathcal{F} be the set of all acts.

Consider an expected utility maximizer with a prior probability μ over S who expects to receive signals from some set Σ . Let $r : S \times \Sigma \rightarrow [0, 1]$ be the information structure he faces, where $r_s(\sigma)$ denotes the probability of learning signal $\sigma \in \Sigma$ given $s \in S$. Let $(\Pr(\cdot | \sigma))_{\sigma \in \Sigma}$ be the corresponding set of posteriors. We say that the pair (μ, r) satisfies the *non-shifting weights property* if (i) $r_s(\sigma) = 0$ implies $\Pr(s | \sigma) = 0$, and (ii) the relative weight of any two states s, s' with $r_s(\sigma) > 0$ and $r_{s'}(\sigma) > 0$ remains intact when σ is learned, that is,

$$\frac{\mu(s)}{\mu(s')} = \frac{\Pr(s | \sigma)}{\Pr(s' | \sigma)}. \quad (1)$$

Consider $f, g \in \mathcal{F}$ and $\sigma \in \Sigma$ such that $f(s) \neq g(s)$ only if $r_s(\sigma) = 0$, and such that $\sum_{s \in S} f(s) \mu(s) \geq \sum_{s \in S} g(s) \mu(s)$. It follows immediately from (1) that for some s' with $r_{s'}(\sigma) > 0$,

$$\begin{aligned} \sum_s f(s) \Pr(s | \sigma) &= \frac{\Pr(s' | \sigma)}{\mu(s')} \sum_{s: r_s(\sigma) > 0} f(s) \mu(s) \\ &\geq \frac{\Pr(s' | \sigma)}{\mu(s')} \sum_{s: r_s(\sigma) > 0} g(s) \mu(s) = \sum_{s \in S} g(s) \Pr(s | \sigma) \end{aligned}$$

In words, the behavioral implication of non-shifting weights is that the ranking of two acts does not change when learning a signal, as long as none of the states where they differ are ruled out. When information is observable, this is what we refer to as *inertial behavior*.

We now establish that non-shifting weights is consistent with Bayesian updating if and only if the information structure r can be described as what we term a *generalized partition*.⁶

Definition 1. A function $\rho : 2^{S'} \rightarrow [0, 1]$ is a *generalized partition* of $S' \subseteq S$ if for any $s \in S'$, ρ_s defined by $\rho_s(I) = \begin{cases} \rho(I) & \text{if } s \in I \\ 0 & \text{if } s \notin I \end{cases}$ satisfies $\sum_{I \subseteq S'} \rho_s(I) = 1$.

A generalized partition induces an information structure where $\rho_s(I)$ is the probability of event I being reported contingent on the state being s . (When there is no danger of

⁶Throughout the paper, \subseteq (resp., \subset) denotes weak (resp., strict) set inclusion. The support of any function is denoted by $\text{supp}(\cdot)$.

confusion, we simply refer to ρ itself as an information structure.) The special case of a set partition corresponds to ρ taking only two values, zero and one. In that case, for every $s \in S'$ there exists a unique $I_s \in 2^{S'}$ with $s \in I_s$ and $\rho_s(I_s) = 1$. Furthermore, $s' \in I_s$ implies that $I_s = I_{s'}$, that is, $\rho_{s'}(I_s) = 1$ for all $s' \in I_s$.

The fact that $\rho_s(I)$ is independent of s (conditional on $s \in I$) implies that even a Bayesian DM cannot draw any inferences from learning an event other than that states outside that event were not realized – in this sense signals are appropriately described as events. Formally,

Proposition 1. *The pair $(\mu, (\Pr(\cdot|I))_{I \in \mathcal{I} \subseteq 2^S})$ satisfies the non-shifting weights property if and only if we can label signals as events, $\Sigma = \mathcal{I}$, and there is a generalized partition ρ such that $r_s \equiv \rho_s$.*

Proof. For generalized partition ρ , Bayes' law implies that for any $s, s' \in I$,

$$\frac{\Pr(s|I)}{\Pr(s'|I)} = \frac{\rho_s(I) \mu(s) / \mu(I)}{\rho_{s'}(I) \mu(s') / \mu(I)} = \frac{\mu(s)}{\mu(s')}$$

independently of I . Conversely, if the information structure r is not a generalized partition, then there exists a signal $\sigma \in \Sigma$ and two states s, s' with $r_s(\sigma) > 0$ and $r_{s'}(\sigma) > 0$, such that $r_s(\sigma) \neq r_{s'}(\sigma)$. For these two states, condition (1) fails. ■

In the next sections we turn to describe how, with unobservable (i.e., subjective) information, inertial behavior can be explained as that of a Bayesian who perceives a generalized partition.

2.2. Inertial Behavior with Unobservable Information

Since we consider an environment where the information structure is subjective and ex post preferences (after a signal has been received) are not observable, we adopt the menu choice approach, with the interpretation that preferences over menus of acts reflect anticipated updating upon receiving information. Let $\mathcal{K}(\mathcal{F})$ be the set of all non-empty compact subsets of \mathcal{F} . Capital letters denote sets, or menus, and small letters denote acts. For example, a typical menu is $F = \{f, g, h, \dots\} \in \mathcal{K}(\mathcal{F})$. Let \succeq be a binary relation over $\mathcal{K}(\mathcal{F})$. The symmetric and asymmetric parts of \succeq are denoted by \sim and \succ , respectively. Let $fIg \in \mathcal{F}$ be the act that agrees with $f \in \mathcal{F}$ on event I and with $g \in \mathcal{F}$ outside I , that is,

$$fIg(s) = \begin{cases} f(s) & s \in I \\ g(s) & \text{otherwise} \end{cases}$$

In order to capture inertial behavior in terms of preferences over menus, we start by defining the class of *small events*. According to Definition 2 below, event $I \in 2^S$ is small if

and only if a bet on I , $\{cI0\}$, cannot be perfectly substituted by the option to bet on any of its strict subsets, $\{cI'0 \mid I' \subset I\}$. Intuitively, event I is small if and only if the DM considers all states in I possible and can foresee learning at least one signal σ that does not rule out any of those states.

Definition 2. *Event $I \subseteq S$ is small if for some $c > 0$,*

$$\{cI0\} \succ \{cI'0 \mid I' \subset I\}.$$

Note that contingent on any signal σ with $r_s(\sigma) = 0$, the act $c(I \setminus \{s\})0$ is as good as $cI0$. Consequently, if for every foreseen signal σ there is at least one state $s \in I$ with $r_s(\sigma) = 0$, then $\{cI0\} \sim \{cI'0 \mid I' \subset I\}$. That is, $\{cI'0 \mid I' \subset I\}$ is a perfect substitute for $\{cI0\}$. Conversely, if the DM considers all states in I possible and can foresee a signal σ with $r_s(\sigma) > 0$ for all $s \in I$,⁷ then contingent on learning σ , the act $cI0$ does strictly better than any act in $\{cI'0 \mid I' \subset I\}$, and thus $\{cI0\} \succ \{cI'0 \mid I' \subset I\}$.

Let $\mathcal{F}_+ := \{f \in \mathcal{F} \mid f(s) > 0 \text{ for all } s \in S\}$ be the collection of acts with strictly positive payoffs in all states. Under very mild assumptions, it is sufficient to impose our main axiom only for acts in \mathcal{F}_+ , since \mathcal{F}_+ is dense in \mathcal{F} with respect to the Euclidean metric (viewing acts as vectors in $[0, 1]^{|S|}$). Confining attention to \mathcal{F}_+ is convenient in writing Definition 3 and Axiom 1 below.

Definition 3. *For a small event I and acts $f, g \in \mathcal{F}_+$ with $f(s) \neq g(s)$ only if $s \in I$, define \succeq_I by*

$$f \succeq_I g \Leftrightarrow \{f, g\} \sim \{f\} \cup \{0 \{s\} g \mid s \in I\}.$$

According to Definition 3, $f \succeq_I g$ describes the property that f is better than g as long as no state in I has been ruled out. To see this, note that $g \in \mathcal{F}_+$ and hence for any signal σ and $s \in I$ such that $r_s(\sigma) > 0$, g is better than $0 \{s\} g$. Therefore, the DM will be indifferent between $\{f, g\}$ and $\{f\} \cup \{0 \{s\} g \mid s \in I\}$ upon learning σ only if he expects to choose f from $\{f, g\}$. At the same time, for any signal σ such that $r_s(\sigma) = 0$ for some $s \in I$, the DM will be indifferent between choosing from $\{f, g\}$ and choosing from $\{f\} \cup \{0 \{s\} g\}$.

Axiom 1 (Inertial behavior). *For any small event I and acts $f, g \in \mathcal{F}_+$ with $f(s) \neq g(s)$ only if $s \in I$,*

$$\{f\} \succeq \{g\} \Rightarrow f \succeq_I g.$$

⁷Formally, all states in I should not be null in the sense of Savage (1954), i.e., for every s it is not the case that $\{f\} \sim \{f \{s\} g\}$ for all f and g . Clearly, if $s \in I$ is null, then $c(I \setminus \{s\})0$ is as good as $cI0$, and hence $\{cI0\} \sim \{cI'0 \mid I' \subset I\}$, that is, $\{cI'0 \mid I' \subset I\}$ is a perfect substitute for $\{cI0\}$.

The axiom captures our interpretation of inertial behavior: the ranking of two acts that differ only on event I does not change as long as no state in I has been ruled out.

It is worth mentioning that Axiom 1 is consistent with the possibility that the DM has strict preference for flexibility, in the sense that it does not preclude the ranking $\{f, g\} \succ \{f\} \succeq \{g\}$. That is, while given I the DM would prefer f to g when no states in I have been ruled out, it may well be that for some σ with $r_s(\sigma) = 0$ for some $s \in I$, g is better than f .

2.3. Generalized Partition Representation

The following utility function over menus incorporates the notion of a generalized partition.

Definition 4. *The pair (μ, ρ) is a generalized-partition representation if (i) $\mu : S \rightarrow [0, 1]$ is a probability measure; (ii) $\rho : 2^{\text{supp}(\mu)} \rightarrow [0, 1]$ is a generalized partition of $\text{supp}(\mu)$; and (iii)*

$$V(F) = \sum_{I \in 2^{\text{supp}(\mu)}} \max_{f \in F} \left[\sum_{s \in I} f(s) \mu(s) \right] \rho(I)$$

represents \succeq .

According to Definition 4, the DM behaves as if he (i) has prior beliefs μ over the state space; and (ii) perceives a generalized partition that determines which event he will learn contingent on the true state. Upon learning event I , the DM calculates posterior beliefs using Bayes' law, which leads him not to shift weights; he excludes all states that are not in I and keeps the relative likelihood of the remaining states fixed. He then chooses from the menu the act that maximizes the corresponding expected utility.

The two parameters of the generalized partition representation, the prior μ on S and the generalized partition ρ , are independent, in the sense that the definition places no joint restriction on the two.⁸ In other words, a generalized-partition representation accommodates any prior beliefs the DM might have about the objective state space, combined with any information structure that can be described as a generalized partition. It is evident from the description of the model that the information structure in Definition 4 is not objectively given. Instead, the generalized-partition should be derived from choice behavior.

Our representation theorem builds on the notion of a subjective-learning representation of \succeq , introduced in DLST

⁸The only role of the restriction that ρ be a generalized partition of $\text{supp}(\mu)$ rather than S is to facilitate the statement of the identification result in Theorem 1: which event is reported contingent on states the DM considers impossible has no behavioral implications, and hence the value of ρ would not be identified on subsets of $S \setminus \text{supp}(\mu)$.

Definition 5. A subjective-learning representation is a function $V : \mathcal{K}(\mathcal{F}) \rightarrow \mathbb{R}$, such that

$$V(F) = \int_{\Delta(S)} \max_{f \in F} \left(\sum_{s \in S} f(s) \pi(s) \right) dp(\pi), \quad (2)$$

where $p(\cdot)$ is a probability measure on $\Delta(S)$, the space of all probability measures on S .

The interpretation of a subjective-learning representation is that the DM behaves as if he has beliefs over the possible posterior distributions that he might face at the time of choosing from the menu. For each posterior $\pi \in \Delta(S)$, he expects to choose from the menu the act that maximizes the corresponding expected utility. As established by DLST (Theorem 1 in their paper), the axioms that are equivalent to the existence of a subjective-learning representation are familiar from the literature on preferences over menus of lotteries – *Ranking*, *vNM Continuity*, *Nontriviality*, and *Independence* – adapted to the domain $\mathcal{K}(\mathcal{F})$, in addition to *Dominance*, which implies monotonicity in payoffs, and *Set Monotonicity*, which captures preference for flexibility. DLST further show that the function $p(\cdot)$ in (2) is unique.

Theorem 1. Suppose that the relation \succeq admits a subjective-learning representation (as in (2)). Then \succeq satisfies Axiom 1 if and only if it has a generalized-partition representation, (μ, ρ) . Furthermore, the pair (μ, ρ) is unique.

Proof. See Appendix B. ■

One immediate implication of a subjective-learning representation (2), is that a forward looking DM with *any* updating rule can be thought of as Bayesian, where the prior that governs the ranking of singletons is simply the expectation over posteriors. Non-shifting weights further requires this prior to agree with *each* posterior on the relative weights of all states that have not been ruled out. In the previous section we verified that this requirement implies that the information structure must be a generalized partition. That is, a forward looking DM who displays inertial behavior (as in Axiom 1) can always be thought of as being rational given some generalized partition.

Compared to the information structures permitted by a subjective-learning representation (as captured by the function $p(\cdot)$ in (2)) a generalized partition (Definition 1) rules out information structures in which two distinct states can generate the same signal but with different probabilities. For example, consider a defendant who is on a trial. Generalized partition is inconsistent with a setting in which there are two states of nature, guilty (G) or innocent (I), and two signals, guilty (g) and innocent (i), such that $\Pr(g|G) > \Pr(g|I) > 0$ and $\Pr(i|I) > \Pr(i|G) > 0$. Obviously, generalized partition is consistent with the

familiar set-partition where signals are deterministic contingent on the state, but it also accommodates many other plausible learning processes. In the next section we provide examples, as well as a full characterization, of those learning processes that give rise to generalized partitions.

In terms of behavior, the following is a simple example of a pattern which our axioms and representation accommodate, but would be precluded if anticipated signals were deterministic. Consider the state space $\{s_1, s_2\}$ and the menu $\{(1, 0), (0, 1), (1 - \varepsilon, 1 - \varepsilon)\}$, which contains the option to bet on either state, as well as an insurance option that pays reasonably well in both states. A DM who is uncertain about the information he will receive by the time he has to choose from the menu may strictly prefer this menu to any of its subsets (for ε small enough). For instance, an investor may value the option to make a risky investment in case he understands the economy well, but also value the option to make a safe investment in case uncertainty remains unresolved at the time of making the investment choice. Our model accommodates this ranking. In contrast, such a ranking of menus is ruled out if signals are deterministic. If the DM expects to learn the true state, then preference for flexibility stems *exclusively* from the DM's prior uncertainty about the true state and the insurance option is irrelevant, that is, $\{(1, 0), (0, 1), (1 - \varepsilon, 1 - \varepsilon)\} \sim \{(1, 0), (0, 1)\}$. And if the DM does not expect to learn the true state, then, for ε small enough, he anticipates choosing the insurance option with certainty, that is, $\{(1, 0), (0, 1), (1 - \varepsilon, 1 - \varepsilon)\} \sim \{(1 - \varepsilon, 1 - \varepsilon)\}$.

3. A Characterization of Generalized Partitions

In this section we characterize the types of learning processes that can give rise to a generalized partition, by identifying all sets of events that support such an information structure. Formally, we characterize the set

$$\left\{ \Psi \subseteq 2^{S'} \mid \text{there is a generalized partition } \rho : 2^{S'} \rightarrow [0, 1] \text{ with } \text{supp}(\rho) = \Psi \right\}.$$

It is worth mentioning that the analysis in this section does not have any decision theoretical aspect and it also does not depend on the distinction between observable and unobservable information.

Definition 6. A collection of events $\Psi \subseteq 2^S$ is a uniform cover of a set $S' \subseteq S$, if (i) $S' = \bigcup_{I \in \Psi} I$; and (ii) there is a function $\beta : \Psi \rightarrow \mathbb{Z}_+$ and a constant $k \geq 1$, such that $\sum_{I \in \Psi | s \in I} \beta(I) = k$ for all $s \in S'$.

To better understand the notion of a uniform cover, consider the following example. Suppose $S = \{s_1, s_2, s_3\}$. Any partition of S , for example $\{\{s_1\}, \{s_2, s_3\}\}$, is a uniform cover of S

(with $k = 1$). A set that consists of multiple partitions, for example $\{\{s_1\}, \{s_2, s_3\}, \{s_1, s_2, s_3\}\}$, is a uniform cover of S (in this example with $k = 2$). The set $\Psi = \{\{s_2, s_3\}, \{s_1, s_2, s_3\}\}$ is *not* a uniform cover of S , because $\sum_{I|s_1 \in I} \beta(I) < \sum_{I|s_2 \in I} \beta(I)$ for any $\beta : \Psi \rightarrow \mathbb{Z}_+$. The set $\{\{s_2, s_3\}, \{s_1\}, \{s_2\}, \{s_3\}\}$, however, is a uniform cover of S with

$$\beta(I) = \begin{cases} 2 & \text{if } I = \{s_1\} \\ 1 & \text{otherwise} \end{cases}.$$

Lastly, the set $\{\{s_1, s_2\}, \{s_2, s_3\}, \{s_1, s_3\}\}$ is a uniform cover of S (with $k = 2$), even though it does not contain a partition.⁹

An empirical situation that gives rise to a uniform cover consisting of two partitions is an experiment that reveals the state of the world if it succeeds, and is completely uninformative otherwise. For a concrete example that gives rise to a uniform cover that does not contain a partition, consider the sequential elimination of n candidates, say during a recruiting process. If k candidates are to be eliminated in the first stage, then the resulting uniform cover is the set of all $(n - k)$ -tuples.

Theorem 2. *A collection of events Ψ is a uniform cover of $S' \subseteq S$ if and only if there is a generalized partition $\rho : 2^{S'} \rightarrow [0, 1]$ with $\text{supp}(\rho) = \Psi$.*

Proof. See Appendix C ■

The ‘only if’ part in the proof of Theorem 2 amounts to finding a solution to a system of linear equations, showing that any uniform cover can give rise to a generalized partition. To illustrate the idea, consider the collection $\{\{s_1\}, \{s_2, s_3\}, \{s_1, s_2, s_3\}\}$ discussed above, which consists of multiple partitions and is a uniform cover with $k = 2$. An information structure r that can be described as a generalized partition

$$\rho : \{\{s_1\}, \{s_2, s_3\}, \{s_1, s_2, s_3\}\} \rightarrow (0, 1]$$

should satisfy two conditions. First, $r_s(I)$ – the probability of receiving the signal I given state s – should be positive and identical for any $s \in I$, that is,

$$\begin{aligned} r_{s_1}(\{s_1\}) &=: \rho(\{s_1\}) > 0 \\ r_{s_2}(\{s_2, s_3\}) &= r_{s_3}(\{s_2, s_3\}) =: \rho(\{s_2, s_3\}) > 0 \\ r_{s_1}(\{s_1, s_2, s_3\}) &= r_{s_2}(\{s_1, s_2, s_3\}) = r_{s_3}(\{s_1, s_2, s_3\}) =: \rho(\{s_1, s_2, s_3\}) > 0 \end{aligned}$$

⁹The notion of uniform cover is closely related to that of balanced collection of weights, which Shapley (1967) introduces in the context of cooperative games.

and 0 otherwise. Second, each $r_s(\cdot)$ must be a probability measure, that is,

$$\begin{aligned} r_{s_1}(\{s_1\}) + r_{s_1}(\{s_1, s_2, s_3\}) &= 1 \\ r_{s_2}(\{s_2, s_3\}) + r_{s_2}(\{s_1, s_2, s_3\}) &= 1 \\ r_{s_3}(\{s_2, s_3\}) + r_{s_3}(\{s_1, s_2, s_3\}) &= 1 \end{aligned}$$

The solutions of this system are given by $\rho(\{s_1, s_2, s_3\}) = \alpha$ and $\rho(\{s_1\}) = \rho(\{s_2, s_3\}) = 1 - \alpha$, for any $\alpha \in (0, 1)$.

Theorem 2 characterizes the types of learning that can be accommodated by a generalized partition. To illustrate it, let us consider a specific example. An oil company is trying to learn whether there is oil in a particular location. Suppose the company can perform a test-drill to determine accurately whether there is oil, $s = 1$, or not, $s = 0$. In that case, the company learns the uniform cover $\Psi = \{\{0\}, \{1\}\}$ that consists of a partition of the state space, and $\rho(\{0\}) = \rho(\{1\}) = 1$ provides a generalized partition. Now suppose that there is a positive probability that the test may not be completed (for some exogenous reason, which is not indicative of whether there is oil or not). The company will either face the trivial partition $\{\{0, 1\}\}$, or the partition $\{\{0\}, \{1\}\}$, and hence $\Psi = \{\{0, 1\}, \{0\}, \{1\}\}$. Suppose the company believes that the experiment will succeed with probability q . Then $\rho(\{0, 1\}) = 1 - q$ and $\rho(\{0\}) = \rho(\{1\}) = q$ provides a generalized partition.

We can extend the previous example and suppose the company is trying to assess the size of an oil field by drilling in l proximate locations, which means that the state space is now $\{0, 1\}^l$. As before, any test may not be completed, independently of the other tests. This is an example of a situation where the state consists of l different attributes (i.e., the state space is a product space), and the DM may learn independently about any of them. Such learning about attributes also gives rise to a uniform cover that consists of multiple partitions and can be accommodated.

To find a generalized partition based on (i) a uniform cover Ψ of a state space S , for which there is a collection Π of partitions whose union is Ψ ; and (ii) a probability distribution q on Π , one can set $\rho(I) = \sum_{\mathcal{P} \in \Pi | I \in \mathcal{P}} q(\mathcal{P})$. We refer to the pair (q, Π) as a random partition.

A different situation in which the DM effectively faces a random partition, (q, Π) , is one where learning is governed by a given filtration of S , but the speed of learning is uncertain. In that case, distinct partitions in Π simply correspond to the information the DM might have at the time of choice, and should thus be ordered by fineness. In that case, $q(\mathcal{P})$ captures the probability of arriving at \mathcal{P} by the time of choice.

Lastly, reconsider the example of sequential elimination of candidates outlined above. Suppose that one out of three candidates will be selected. Write the corresponding state space as $S = \{s_1, s_2, s_3\}$. If one candidate will be eliminated in the first round, then the

uniform cover of events the DM might learn is given by $\Psi = \{\{s_1, s_2\}, \{s_2, s_3\}, \{s_1, s_3\}\}$. Suppose that, contingent on person i being the best candidate, the DM considers any order of elimination of the other candidates as equally likely. This corresponds to the generalized partition with $\rho(I) = 0.5$ for all $I \in \Psi$ and $\rho(I) = 0$ otherwise.

Each of the generalized partitions discussed in this section can be coupled with any prior beliefs μ to generate a generalized-partition representation (μ, ρ) .

4. Other Interpretations of Generalized Partitions

So far we have assumed the analyst is completely unaware of the underlying information structure. However, in many applications, the analyst will have at least some knowledge of the true information structure, which may be inconsistent with the generalized partition that explains the observed behavior according to Theorem 1. To reconcile this gap, we now investigate two different protocols that may lead the DM to subjectively perceive a generalized partition (and to display inertial behavior) even if the true information structure is not.

The first protocol is based on the psychological notion of cognitive inertia. The second protocol is inspired by a concept from social psychology, called construal level theory, which contends that the closer in time a situation is, the more easily and concretely it will be thought of. Applied to our context, we study a boundedly rational agent who is constrained in his ability to directly assess probabilities for the outcome of random variables that resolve in the distant future, and tries to behave rationally given this constraint.

4.1. Cognitive Inertia

Behavioral psychologists use the notion of cognitive inertia to describe the “tendency for beliefs or sets of beliefs to endure once formed”.¹⁰ For instance, one explanation for the difficulties of implementing organizational change is that managers fail to update and revise their understanding of a situation even when the environment changes, say when new technologies emerge (Tripsas and Gavett 2000). This impedes their ability to adjust their actions in response to relevant information, and may lead to inertial behavior.

Theorem 1 does not preclude the interpretation that our DM ‘suffers’ from cognitive inertia, but implies that it is restrictive. In our formal environment, we interpret cognitive inertia to mean that the only information the DM reacts to is which states have conclusively been ruled out, and that this information is taken into account following Bayes’ law. Our result then establishes that, for a forward looking DM (for whom the expectation over

¹⁰https://en.wikipedia.org/wiki/Cognitive_inertia

posteriors coincides with the prior), cognitive inertia can explain behavior only if the *true* information structure is such that the *biased* one — which is derived from the true one by ignoring all information about relative weights — is a generalized partition. In this case it is unnecessary to model the DM as biased, even if he is subject to cognitive inertia; learning by generalized partition provides a rational *as if* model of his behavior.

4.2. Shortsighted Rationality

This section considers a DM who receives a partial (but accurate) description of the true information structure, namely for each event the probability that it will be ruled out in the first stage, and who interprets these probabilities as being generated by some generalized partition as well as a probability distribution over the state space.

While the constraint on the DM’s knowledge could be exogenous, we will interpret it as a bound on the horizon up to which the DM can independently form beliefs about his uncertain environment. We call such a DM *shortsighted*. Specifically, in our model uncertainty about the payoff relevant state resolves in two stages, and we assume that the shortsighted DM can only directly assign probabilities to uncertainty that resolves in the first stage. This constraint has two consequences. First, the DM does not initially hold beliefs over the final states, but only over the signals he might learn in the first period. Second, he is oblivious to possible correlation in the resolution of uncertainty across periods. In particular, he cannot tell apart two signals with the same support that induce different posterior probabilities over the states (see, for example the defendant example in Section 2.3). Since the DM has no perception of correlation, it is without loss to treat those two signals as the same and to identify each signal with the event consisting of those states that are not ruled out by it.

Being otherwise rational, the shortsighted DM would like to infer the distribution over states, as well as the underlying information structure, from his beliefs over first period events. Since he does not perceive any correlation between events and states (beyond his understanding of which states have been ruled out), it seems natural that he will consider only information structures that likewise feature no correlation, that is, he will try to find a generalized partition. We call this protocol for inference *shortsightedly rational*, and ask whether following the protocol is always possible. The next result answers this question in the affirmative: it establishes that the DM can always find a pair of prior and generalized partition that jointly generate the given distribution over events.¹¹

¹¹Once more, Theorem 3 and Theorem 4 below are not restricted to the analysis of a boundedly rational DM, but can be more broadly applied whenever the DM receives a partial description of the true information structure, in the form of the probability distribution over events it induces.

Theorem 3. *Let π be a probability measure over $2^S \setminus \emptyset$. There exists $\mu \in \Delta(S)$ and a generalized partition ρ on $2^{\sigma(\mu)}$ such that $\rho(I)\mu(I) = \pi(I)$ for all $I \in 2^{\sigma(\mu)}$.*

A proof is in Appendix D. Observe that such μ and ρ exist if and only if for any $s \in S$ either $\mu(s) = 0$ or $\sum_{I \in 2^{\sigma(\mu)} | s \in I} \rho(I) = \sum_{I \in 2^{\sigma(\mu)} | s \in I} \frac{\pi(I)}{\mu(I)} = 1$ holds. We establish that a solution to this non-linear system of equations exists via a fixed point theorem. In light of Theorem 2, it is clear that the support of μ must be such that the support of π constitutes a uniform cover of it.

The pair (μ, ρ) in the theorem may not be unique: Suppose, for example, that for all $I \in \text{supp}(\pi)$ it is the case that $s \in I$ if and only if $s' \in I$. Then the relative weight μ assigns to s versus s' can not be identified. However, as the following result shows, uniqueness is guaranteed if the DM can conceive of learning the state already in the first stage (i.e., $\pi(\{s\}) > 0$ for all $s \in S$).

Theorem 4. *Let π be a probability measure over $2^S \setminus \emptyset$ with $\pi(\{s\}) > 0$ for all $s \in S$. If $\mu \in \Delta(S)$ and ρ is a generalized partition on $2^{\sigma(\mu)}$ such that $\rho(I)\mu(I) = \pi(I)$ for all $I \in 2^{\sigma(\mu)}$, then for any other (μ', ρ') with this feature, $\mu' = \mu$ and $\rho' = \rho$.*

A proof is in Appendix E. To recapitulate, Theorem 3 implies that a shortsighted DM who follows the protocol described above will display cognitive inertia for *any* information structure he faces. Suppose the analyst happens to know the true information structure and that the condition in Theorem 4 is met. If the analyst is also willing to assume that the DM is shortsightedly rational, then the theorem implies that he can determine the information structure inferred by the DM, without relying on any choice data.

5. Related Literature

Our model builds on a general subjective-learning representation of preferences over menus of acts, which identifies signals with posterior distributions over states (Definition 5 in Section 2.3). Such representation was first derived in DLST. They further derived a partitional learning representation, in which signals correspond to events that partition the state space. Partitional learning is also studied in De Oliveira et al. (2016), who adopt the same menu choice approach but, unlike us, allow the DM to choose his information structure, and in Lu (2015), who uses the random choice approach to elicit information. A dynamic version of partitional learning appears in Dillenberger, Krishna, and Sadowski (2016).

In terms of generality, the generalized partition representation lies in between the two extreme cases studied in DLST. For a Bayesian DM, a salient feature of partitional learning

is that the relative weights of states in the learned event do not shift. In this paper we show that this non-shifting-weights property corresponds to inertial behavior, which links ex ante and ex post rankings of acts. We then identify the *least restrictive* model with this property. Sections 2.3 and 3 demonstrate the merit of the generality permitted by our model (over partitional learning). In Appendix A we show how one can strengthen our main axiom to give a different characterization of partitional learning than the one provided in the aforementioned papers.

Starting with Sims (1998, 2003), numerous models of rational inattention were suggested to explain and match the inertia found in macroeconomic data, such as sluggish adjustment of prices in reaction to demand shocks (see, among many others, Mackowiak and Wiederholt 2009 and 2015; Steiner, Stewart, and Matějka 2016; and the empirical findings in Andersen et al. 2015). These models are based on the assumption that economic agents have limited attention which they allocate optimally, leading them to update their beliefs only infrequently. The inertia addressed by these models is, then, a delay in the reaction to new information. In contrast, inertial behavior in our paper captures immediate but potentially incomplete adjustment. Another related model, though one without any information arrival, is proposed in Vega-Redondo (1993). In his model, risk averse agents have precise perceptions of the consequences associated with the action they are currently performing, but only imprecise perceptions of all other actions, leading to status quo bias (which he refers to as inertial behavior).

Our notion of inertial behavior resembles the following familiar requirement, which is often referred to as dynamic consistency in the context of objective information: the ex ante comparison of alternatives that differ only on some event $I \subseteq S$ is the same as the ex post comparison of those alternatives, contingent upon learning that event I has occurred.¹² Breaking this requirement into two parts, violations of the axiom might emerge from either (i) an inconsistency between ex ante preferences and anticipated ex post preferences, or (ii) an inconsistency between anticipated ex post preferences and actual ex post preferences. Siniscalchi (2011), for example, explicitly focusses on (ii).¹³ Importantly, in Siniscalchi's paper the information structure is objective – modeled as a fixed filtration, or partition if restricted to two stages – and signals are observable, so that it is possible to elicit anticipated ex post preferences by directly asking the DM whether he prefers to receive act f or g contingent on learning event I , while otherwise receiving the same act h in either case. In contrast, in our paper information is subjective, and hence inertial behavior rules out (i), but must be silent on (ii).

¹²See, for example, Axiom 2 in Ghirardato (2002).

¹³Siniscalchi refers to anticipated ex post preferences as *conjectural* ex post preferences.

Appendices

A. Comparison with Partitional Learning

We now show how Axiom 1 can be modified to capture the idea of deterministic signals and thereby to characterize partitional learning. The axiomatization here is different than the one provided in DLST.

Axiom 2 (Partition). *For any two acts $f, g \in \mathcal{F}_+$ and small event $I \subseteq S$,*

$$\{fI0\} \succeq \{gI0\} \Rightarrow \{fI0, gI0\} \sim \{fI0\}.$$

In the words of Kreps (1979), Axiom 2 can be viewed as a requirement of strategic rationality given an event. Recall that if I is small then the DM can foresee learning an event J that (weakly) contains I . If signals are deterministic and I is small, then there is no $I' \subset I$ the DM also foresees learning. It is thus obvious why for a Bayesian DM partitional learning satisfies this property; $fI0$ and $gI0$ agree outside I , and if $fI0$ does better on I , it will do better on any superset of I . Conversely, if information in the subjective-learning representation of Definition 5 is not partitional, then there are two posteriors $\pi, \pi' \in \text{supp}(p)$, such that $\text{supp}(\pi) = I$, $\text{supp}(\pi') = I'$, $I \neq I'$, and $I \cap I' \neq \emptyset$. Let $f = c$ and $g = c + \varepsilon$. Without loss of generality, suppose $I \cap I' \neq I$. If $\varepsilon > 0$ and small enough, then $\{fI0\} \succeq \{g(I \cap I')0\}$ while, since $g(I \cap I')0$ does better than $fI0$ conditional on I' , $\{fI0, g(I \cap I')0\} \succ \{fI0\}$.

Since larger sets are better (preference for flexibility), $\{fI0\} \cup \{gI'0 | I' \subset I\} \succeq \{fI0\}$. Further, by payoff domination, $\{fI0, gI0\} \succeq \{fI0\} \cup \{gI'0 | I' \subset I\}$. Therefore, Axiom 2 implies Axiom 1, while, as is obvious from our results, the converse is false.

B. Proof of Theorem 1

Necessity: Suppose \succeq admits a generalized partition representation (μ, ρ) , that is,

$$V(F) = \sum_{I \in 2^{\text{supp}(\mu)}} \max_{f \in F} \left[\sum_{s \in I} f(s) \mu(s) \right] \rho(I)$$

represents \succeq .

Claim 1. *Event $I \subseteq S$ is small if and only if $\rho(J) > 0$ for some $J \supseteq I$.*

Proof. This follows immediately from the following two observations:

1. For any J with $J \cap I \neq I$,

$$\max_{f \in \{cI'0 | I' \subset I\}} [\sum_{s \in J} f(s) \mu(s)] = c \sum_{s \in J \cap I} \mu(s) = \sum_{s \in J} cI0(s) \mu(s);$$

2. For any $J \supseteq I$, $\max_{f \in \{cI'0 | I' \subset I\}} [\sum_{s \in J} f(s) \mu(s)] < c \sum_{s \in I} \mu(s) = \sum_{s \in J} cI0(s) \mu(s)$.

■

Consider a small event I and two acts $f, g \in \mathcal{F}_+$. According to the generalized partition representation,

$$\{fI0\} \succeq \{gI0\} \Leftrightarrow \sum_{s \in I} [f(s) \mu(s) \sum_{I \subseteq S} \rho_s(I)] \geq \sum_{s \in I} [g(s) \mu(s) \sum_{I \subseteq S} \rho_s(I)]$$

Since ρ is a generalized partition, $\sum_{I \subseteq S} \rho_s(I) = 1$ for all $s \in \text{supp}(\mu)$, and hence $\{fI0\} \succeq \{gI0\} \Leftrightarrow \sum_{s \in I} f(s) \mu(s) \geq \sum_{s \in I} g(s) \mu(s)$. Therefore, $\{fI0\} \succeq \{gI0\}$ implies that for any event $J \supseteq I$ with $J \in \text{supp}(\rho)$ we have

$$\begin{aligned} \max_{f' \in \{fI0, gI0\}} [\sum_{s \in J} f'(s) \mu(s)] &= \max [\sum_{s \in I} f(s) \mu(s), \sum_{s \in I} g(s) \mu(s)] = \\ &= \sum_{s \in I} f(s) \mu(s) = \max_{f' \in \{fI0\} \cup \{gI'0 | I' \subset I\}} [\sum_{s \in J} f'(s) \mu(s)] \end{aligned}$$

and also implies that for any event $J \in \text{supp}(\rho)$ with $J \cap I \neq I$ we have

$$\max_{f' \in \{fI0, gI0\}} [\sum_{s \in J} f'(s) \mu(s)] = \max_{f' \in \{fI0\} \cup \{gI'0 | I' \subset I\}} [\sum_{s \in J} f'(s) \mu(s)].$$

Therefore, $\{fI0\} \succeq \{gI0\} \Rightarrow \{fI0, gI0\} \sim \{fI0\} \cup \{gI'0 | I' \subset I\}$.

Sufficiency: Suppose that \succeq admits a subjective learning representation and satisfies Axiom 1. For each $s \in S$, let $\mu(s) = \int_{\text{supp}(p)} \pi(s) dp(\pi)$. Since the measure p over $\Delta(S)$ in a subjective-learning representation is unique, so is μ . It is easy to see that μ is a probability measure and that for each $g, f \in \mathcal{F}$, $\{f\} \succ \{g\} \Leftrightarrow \sum_{s \in S} f(s) \mu(s) > \sum_{s \in S} g(s) \mu(s)$.

Let:

$$\Pi_0 := \left\{ \pi \left| \exists s, s' \in \text{supp}(\pi) \cap \text{supp}(\mu) \text{ with } \frac{\pi(s)}{\pi(s')} \neq \frac{\mu(s)}{\mu(s')} \right. \right\}$$

and for $\varepsilon > 0$

$$\Pi_\varepsilon := \left\{ \pi \left| \exists s, s' \in \text{supp}(\pi) \cap \text{supp}(\mu) \text{ with } \frac{\pi(s)}{\pi(s')} \geq (1 + \varepsilon) \frac{\mu(s)}{\mu(s')} \right. \right\}.$$

Finally, for any $s, s' \in S$ let

$$\Pi_{s,s',\varepsilon} := \left\{ \pi \mid s, s' \in \text{supp}(\pi) \cap \text{supp}(\mu) \text{ and } \frac{\pi(s)}{\pi(s')} \geq (1 + \varepsilon) \frac{\mu(s)}{\mu(s')} \right\}.$$

Claim 2. *If $p(\Pi_0) > 0$, then there exist $s, s' \in \text{supp}(\mu)$ and $\varepsilon > 0$, such that $p(\Pi_{s,s',\varepsilon}) > 0$.*

Proof. First note that there exists $\varepsilon > 0$ such that $p(\Pi_\varepsilon) > 0$. If it was not the case, then there would be $(\varepsilon_n) \downarrow 0$ such that

$$p\left(\left\{ \pi \mid \exists s, s' \in \text{supp}(\pi) \cap \text{supp}(\mu) \text{ with } \frac{\pi(s)}{\pi(s')} \notin \left((1 - \varepsilon_n) \frac{\mu(s)}{\mu(s')}, (1 + \varepsilon_n) \frac{\mu(s)}{\mu(s')} \right) \right\}\right) = 0$$

for all n . But for any s, s' we would have $\left((1 - \varepsilon_n) \frac{\mu(s)}{\mu(s')}, (1 + \varepsilon_n) \frac{\mu(s)}{\mu(s')} \right) \rightarrow \frac{\mu(s)}{\mu(s')}$ and therefore $p(\Pi_0) = 0$, which is a contradiction.

Second, since S is finite (and thereby also $\text{supp}(\mu)$, and $\text{supp}(\mu) \times \text{supp}(\mu)$) we can write $p(\Pi_\varepsilon) = \sum_{s,s' \in \text{supp}(\mu) \times \text{supp}(\mu)} p(\Pi_{s,s',\varepsilon})$. This immediately implies that at least for one pair s, s' we must have $p(\Pi_{s,s',\varepsilon}) > 0$. ■

Now fix s, s' such that $p(\Pi_{s,s',\varepsilon}) > 0$. Let $f := \frac{1}{2}$, and for $\gamma, \beta > 0$ such that $\frac{\beta}{\gamma} \in \left(\frac{\mu(s)}{\mu(s')}, (1 + \varepsilon) \frac{\mu(s)}{\mu(s')} \right)$, define an act g by

$$g(\hat{s}) := \begin{cases} \frac{1}{2} + \gamma & \hat{s} = s \\ \frac{1}{2} - \beta & \hat{s} = s' \\ \frac{1}{2} & \text{otherwise} \end{cases}$$

Note that for any I with $s, s' \in I$

$$\pi \in \Pi_{s,s',\varepsilon} \Rightarrow \frac{1}{2} < \sum_{\hat{s} \in I} g(\hat{s}) \pi(\hat{s}); \text{ and} \quad (3i)$$

$$\frac{1}{2} > \sum_{\hat{s} \in I} g(\hat{s}) \mu(\hat{s}). \quad (3ii)$$

Let $\mu_{s,s',\varepsilon}(s) := \int_{\Pi_{s,s',\varepsilon}} \pi(s) dp(\pi)$. We now argue that there exists a small $I \subseteq \text{supp}(\mu_{s,s',\varepsilon})$ such that

(a) $f \not\prec_I g$;

(b) $\{fI0\} \succeq \{gI0\}$.

First, $s, s' \in \text{supp}(\pi)$ for every $\pi \in \Pi_{s,s',\varepsilon}$. Because S is finite, $\Pi_{s,s',\varepsilon}$ can be partitioned into finitely many sets of posteriors, where π and π' are in the same cell of the partition if

and only if they have the same support. At least one of these cells must have positive weight under p . Choose I to be such a cell. By (3i), we have $\{fI0, gI0\} \succ \{fI0\} \cup \{gI'0 \mid I' \subset I\}$, or $f \not\prec_I g$. This establishes part (a). Part (b) is immediately implied by (3ii).

The combination of (a) and (b) constitutes a violation of Axiom 1. Therefore $p(\Pi_0) = 0$, that is, for all but a measure zero of posteriors, $\pi \in \text{supp}(p)$ and $I \subset S$ imply that

$$\pi(s) = \begin{cases} \frac{\mu(s)}{\mu(I)} & s \in I \\ 0 & \text{otherwise} \end{cases}$$

Note that this result implies that if $\pi, \pi' \in \text{supp}(p)$ are such that $\pi \neq \pi'$ then $\text{supp}(\pi) \neq \text{supp}(\pi')$. Therefore, we can index each element $\pi \in \text{supp}(p)$ by its support $\text{supp}(\pi) \in 2^{\text{supp}(\mu)}$ and denote a typical element by $\pi(\cdot \mid I)$, where $\pi(s \mid I) = 0$ if $s \notin I \in 2^{\text{supp}(\mu)}$. There is then a unique measure \hat{p} such that for all $F \in \mathcal{K}(\mathcal{F})$

$$V(F) = \sum_{I \in 2^{\text{supp}(\mu)}} \max_{f \in F} \left[\sum_{s \in S} f(s) \pi(s \mid I) \right] \hat{p}(I), \quad (3)$$

and $\mu(s) = \sum_{I \mid s \in I} \pi(s \mid I) \hat{p}(I)$.

We have already established that $\pi(s \mid I) = \frac{\mu(s)}{\mu(I)}$ for all $s \in I \in \text{supp}(\hat{p})$. Define $\rho(I) := \frac{\hat{p}(I)}{\mu(I)}$ and substitute $\mu(s) \rho(I)$ for $\pi(s \mid I) \hat{p}(I)$ in (3). Bayes' law implies that

$$\rho_s(I) := \begin{cases} \rho(I) & \text{if } s \in I \\ 0 & \text{if } s \notin I \end{cases}$$

is indeed a probability measure for all s .

C. Proof of Theorem 2

(if) Let Ψ be a uniform cover of S' . Let $k \geq 1$ be the smallest number of times that S' is covered by Ψ . Set $\rho(I) = \frac{\beta(I)}{k}$ for all $I \in \Psi$.

(only if) Suppose that $\rho : 2^{S'} \rightarrow [0, 1]$ is a generalized partition, with $\text{supp}(\rho) = \Psi$. In addition to $\rho(I) = 0$ for $I \notin \Psi$, the conditions that ρ should satisfy can be written as $\mathbf{A}\rho_\Psi = \mathbf{1}$, where \mathbf{A} is a $|S'| \times |\Psi|$ matrix with entries $a_{i,j} = \begin{cases} 1 & s \in I \\ 0 & s \notin I \end{cases}$, ρ_Ψ is a $|\Psi|$ -dimensional vector with entries $(\rho(I))_{I \in \Psi}$, and $\mathbf{1}$ is a $|S'|$ -dimensional vector of ones.

Suppose first that $\rho(I) \in \mathbb{Q} \cap (0, 1]$ for all $I \in \Psi$. Rewrite the vector ρ_Ψ by expressing all entries using the smallest common denominator, $\xi \in \mathbb{N}_+$. Then Ψ is a generalized partition of size ξ . To see this, let $\beta(I) := \xi \rho(I)$ for all $I \in \Psi$. Then $\sum_{I \in \Psi \mid s \in I} \beta(I) = \sum_{I \in \Psi \mid s \in I} \xi \rho(I) = \xi$ for all $s \in S'$.

It is thus left to show that if $\rho_\Psi \in (0, 1]^{|\Psi|}$ solves $\mathbf{A}\rho_\Psi = \mathbf{1}$, then there is also $\rho'_\Psi \in [\mathbb{Q} \cap (0, 1)]^{|\Psi|}$ such that $\mathbf{A}\rho'_\Psi = \mathbf{1}$.

Let \widehat{P} be the set of solutions for the system $\mathbf{A}\rho_\Psi = \mathbf{1}$. Then, there exists $X \in \mathbb{R}^k$ (with $k \leq |\Psi|$) and an affine function $f : X \rightarrow \mathbb{R}^{|\Psi|}$ such that $\widehat{\rho}_\Psi \in \widehat{P}$ implies $\widehat{\rho}_\Psi = f(x)$ for some $x \in X$. We first make the following two observations:

- (i) there exists f as above, such that $x \in \mathbb{Q}^k$ implies $f(x) \in \mathbb{Q}^{|\Psi|}$;
- (ii) there exists an open set $\widetilde{X} \subseteq \mathbb{R}^k$ such that $f(x) \in \widehat{P}$ for all $x \in \widetilde{X}$

To show (i), apply the Gauss elimination procedure to get f and X as above. Using the assumption that \mathbf{A} has only rational entries, the Gauss elimination procedure (which involves a sequence of elementary operations on \mathbf{A}) guarantees that $x \in \mathbb{Q}^k$ implies $f(x) \in \mathbb{Q}^{|\Psi|}$.

To show (ii), suppose first that $\rho^* \in \widehat{P} \cap (0, 1)^{|\Psi|}$ and $\rho^*_\Psi \notin \mathbb{Q}^{|\Psi|}$. By construction, $\rho^*_\Psi = f(x^*)$, for some $x^* \in X$. Since $\rho^*_\Psi \in (0, 1)^{|\Psi|}$ and f is affine, there exists an open ball $B_\varepsilon(x^*) \subset \mathbb{R}^k$ such that $f(x) \in \widehat{P} \cap (0, 1)^{|\Psi|}$ for all $x \in B_\varepsilon(x^*)$, and in particular for $x' \in B_\varepsilon(x^*) \cap \mathbb{Q}^k$ ($\neq \phi$). Then $\rho'_\Psi = f(x') \in [\mathbb{Q} \cap (0, 1)]^{|\Psi|}$. Lastly, suppose that $\rho^*_\Psi \in \widehat{P} \cap (0, 1)^{|\Psi|}$ and that there are $0 \leq l \leq |\Psi|$ sets $I \in \Psi$, for which $\rho(I)$ is uniquely determined to be 1. Then set those l values to 1 and repeat the above procedure for the remaining system of $|\Psi| - l$ linear equations.

D. Proof of Theorem 3

First observe that such μ and ρ exist if, and only if, for any $s \in S$ either $\mu(s) = 0$ or $\sum_{\{I \in 2^{\sigma(\mu)} | s \in I\}} \rho(I) = \sum_{\{I \in 2^{\sigma(\mu)} | s \in I\}} \frac{\pi(I)}{\mu(I)} = 1$ holds. If $\mu(s) > 0$, then $\mu(I) \geq \pi(I)$ for any set $I \in 2^S$. Consequently, let

$$\mathcal{L} = \{\mu \in \Delta(S) \mid \mu(I) \geq \pi(I) \text{ for all } I \in 2^S\}$$

Claim 3. *The set \mathcal{L} is nonempty, closed, and convex.*

Proof of Claim 3: To see that \mathcal{L} is nonempty, consider the measure μ over S defined by $\mu(s_i) = \sum_{\{I | s_i \in I \text{ and } s_j \notin I \text{ for } j < i\}} \pi(I)$. Clearly $\mu(s_i) \in [0, 1]$ for all i and $\sum_i \mu_i = \sum_{\{I \in 2^S \setminus \emptyset\}} \pi(I) = 1$. Furthermore, if we let $i_I = \min\{i | s_i \in I\}$, then $\pi(I) = \mu(s_{i_I})$ and thus $\mu(I) = \sum_{s \in I} \mu(s) \geq \pi(I)$, so $\mu \in \mathcal{L}$. For convexity, take $\mu, \mu' \in \mathcal{L}$. Since $\Delta(S)$ is a mixture space, $\alpha\mu + (1 - \alpha)\mu' \in \Delta(S)$. And since $\mu(I) \geq \pi(I)$ and $\mu'(I) \geq \pi(I)$ for all $I \in 2^S$, $(\alpha\mu + (1 - \alpha)\mu')(I) = \alpha\mu(I) + (1 - \alpha)\mu'(I) \geq \pi(I)$ and hence $\alpha\mu + (1 - \alpha)\mu' \in \mathcal{L}$. Lastly, \mathcal{L} is closed since it is the intersection of closed half-spaces with the closed set $\Delta(S)$. \square

Define $f : \mathcal{L} \rightarrow \mathbb{R}^{|S|}$ by

$$[f\mu](s) = \begin{cases} 0 & \text{if } \mu(s) = 0 \\ \mu(s) \times \left[\sum_{\{I \in \text{supp}(\pi) | s \in I\}} \frac{\pi(I)}{\mu(I)} \right] & \text{otherwise} \end{cases}$$

Claim 4. $f(\mu) \in \mathcal{L}$ and f is continuous.

Proof of Claim 4: To see that $f(\mu) \in \mathcal{L}$, pick $\mu \in \mathcal{L}$. We first show that $f(\mu) \in \Delta(S)$.

To see this, note that $[f\mu](s) \geq 0$ and that

$$\begin{aligned} \sum_{s \in S} [f\mu](s) &= \sum_{s \in S} \mu(s) \times \left[\sum_{\{I \in \text{supp}(\pi) | s \in I\}} \frac{\pi(I)}{\mu(I)} \right] \\ &= \sum_I \sum_{s \in I} \mu(s) \times \left[\frac{\pi(I)}{\mu(I)} \right] \\ &= \sum_I \left[\frac{\pi(I)}{\mu(I)} \right] \sum_{s \in I} \mu(s) \\ &= \sum_I \left[\frac{\pi(I)}{\mu(I)} \right] \mu(I) = \sum_I \pi(I) = 1. \end{aligned}$$

Now pick $I' \in \text{supp}(\pi)$ and note that

$$\sum_{s \in I'} [f\mu](s) = \sum_{s \in I'} \mu(s) \times \left[\sum_{\{I \in \text{supp}(\pi) | s \in I\}} \frac{\pi(I)}{\mu(I)} \right] \geq \sum_{s \in I'} \mu(s) \times \frac{\pi(I')}{\mu(I')} = \pi(I')$$

hence $f(\mu) \in \mathcal{L}$.

For continuity, observe that f is obviously continuous on $\text{int}(\mathcal{L})$. Also notice that $\pi(\{s\}) > 0$ implies $\mu(s) > 0$ for all $\mu \in \mathcal{L}$. Let $S' = \{s \in S | \pi(\{s\}) = 0\} \subset S$. Let $\langle \mu_k \rangle \in \text{int}(\mathcal{L})$ be such that $\mu_k \rightarrow \bar{\mu}$, where $\bar{\mu}(s) = 0$ if, and only if $s \in J$, for some $J \subseteq S$. For any $s \in J$ we have

$$\begin{aligned} [f\mu_k](s) &= \mu_k(s) \times \left[\sum_{\{I \in \text{supp}(\pi) | s \in I\}} \frac{\pi(I)}{\mu_k(I)} \right] \\ &= \mu_k(s) \times \left[\sum_{\{I \in \text{supp}(\pi), I \subseteq S' | s \in I\}} \frac{\pi(I)}{\mu_k(I)} + \sum_{\{I \in \text{supp}(\pi), I \not\subseteq S' | s \in I\}} \frac{\pi(I)}{\mu_k(I)} \right] \\ &= \mu_k(s) \times \left[0 + \sum_{\{I \in \text{supp}(\pi), I \not\subseteq S' | s \in I\}} \frac{\pi(I)}{\mu_k(I)} \right] \end{aligned}$$

and

$$\lim_{\mu(s) \rightarrow 0} \mu(s) \times \sum_{\{I \in \text{supp}(\pi), I \not\subseteq S' | s \in I\}} \frac{\pi(I)}{\mu(I)} = 0.$$

Clearly $[f\mu_k](s) \rightarrow [f\bar{\mu}](s)$ for any $s \notin J$. Therefore, f is continuous. \square

From the two claims above, we conclude that f is a continuous function from a convex compact subset of $\mathbb{R}^{|S|}$ to itself. By Brouwer's fixed-point theorem, f has a fixed point, μ^* . But note that if μ^* is a fixed point of f and $\mu^*(s) > 0$, then $\sum_{\{I \in 2^S \setminus \emptyset \mid s \in I\}} \frac{\pi(I)}{\mu^*(I)} = 1$, and thus the pair (μ^*, ρ) with $\rho(I) = \frac{\pi(I)}{\mu^*(I)}$ satisfies all the requirements. ■

E. Proof of Theorem 4

Suppose, contrary to the theorem, that $\mu \neq \mu'$ (otherwise $\rho = \rho'$ follows immediately.) Without loss, order states such that for some l and j with $k - j > l$ we have $\mu'(s_i) > \mu(s_i)$ if and only if $i \in \{1, \dots, l\}$, and $\mu(s_i) > \mu'(s_i)$ if and only if $i \in \{k - j, \dots, k\}$.

Let $\mu_0 = \mu$, $\bar{s}^0 = s_1$, and $\underline{s}^0 = s_k$. Define inductively $\rho_t(I) := \frac{\pi(I)}{\mu_t(I)}$, $\Delta\mu_t(s) = \mu_t(s) - \mu'(s)$,

$$\mu_{t+1}(s) := \begin{cases} \mu_t(s) & s \notin \{\bar{s}^t, \underline{s}^t\} \\ \mu_t(s) + \min(|\Delta\mu_t(\bar{s}^t)|, |\Delta\mu_t(\underline{s}^t)|) & s = \bar{s}^t \\ \mu_t(s) - \min(|\Delta\mu_t(\bar{s}^t)|, |\Delta\mu_t(\underline{s}^t)|) & s = \underline{s}^t \end{cases}$$

$\bar{s}^t := \arg \min_{s: \mu_t(s) < \mu'(s)} \sum_{I: s \in I} \rho_t(I)$, and $\underline{s}^t := \arg \max_{s: \mu_t(s) > \mu'(s)} \sum_{I: s \in I} \rho_t(I)$. Note that there is finite T with $\mu_T = \mu'$, at which point the process terminates. Hence, if we establish that ρ_t is never a generalized partition for $t > 0$, we are done.

Claim 5 (Induction beginning). ρ_1 is not a generalized partition. In particular

1. $\sum_{I: s_1 \in I} \rho_1(I) < \min \left\{ 1, \min_{s \neq s_1} \left\{ \sum_{I: s \in I} \rho_1(I) \right\} \right\}$
2. $\sum_{I: s_k \in I} \rho_1(I) > \max \left\{ 1, \max_{s \neq s_k} \left\{ \sum_{I: s \in I} \rho_1(I) \right\} \right\}$.

Proof of Claim 5: $\mu_1(I) \geq \mu(I)$ if and only if $s_1 \in I$, with strict inequality for $s_2 \notin I$. Part 1 then follows immediately, and part 2 is analogous. □

Induction hypothesis. ρ_t is not a generalized partition. In particular

1. There is s^t with $\mu(s^t) > \mu_t(s^t)$ and $1 < \sum_{I: s^t \in I} \rho_t(I) \geq \max_{s \neq s^t} \left\{ \sum_{I: s \in I} \rho_t(I) \right\}$.
2. There is s_t with $\mu(s_t) < \mu_t(s_t)$ and $1 > \sum_{I: s_t \in I} \rho_t(I) \leq \min_{s \neq s_t} \left\{ \sum_{I: s \in I} \rho_t(I) \right\}$.

Claim 6 (Induction step). If μ_t and ρ_t satisfy the hypothesis, then so do μ_{t+1} and ρ_{t+1} .

Proof of Claim 6: Note that (i) $\mu_{t+1}(I) > \mu_t(I)$ if and only if $\bar{s}^t \in I$ and $\underline{s}^t \notin I$ (and in particular for $\{\bar{s}^t\}$) and (ii) $\mu_{t+1}(I) < \mu_t(I)$ if and only if $\underline{s}^t \in I$ and $\bar{s}^t \notin I$.

Suppose, contrary to the claim, that ρ_{t+1} is a generalized partition. Then it must be that $\sum_{I:\bar{s}^t \in I} \rho_{t+1}(I) = 1$. But by (1) of the hypothesis, $1 < \sum_{I:\bar{s}^t \in I} \rho_t(I) \geq \sum_{I:\bar{s}^t \in I} \rho_t(I)$ and then by (i) $\sum_{I:\bar{s}^t \in I} \rho_{t+1}(I) < 1$. Hence, ρ_{t+1} is not a generalized partition. In particular, due to (i) and the choice of \bar{s}^t , $\sum_{I:\bar{s}^t \in I} \rho_{t+1}(I) < \sum_{I:s \in I} \rho_{t+1}(I)$ for all s with $\mu_{t+1}(s) < \mu'(s)$, and hence (1) of the claim holds. (2) follows analogously. \square

This establishes the Theorem. \blacksquare

References

- [1] Andersen, S., J. Y. Campbell, K. M. Nielsen, and T. Ramadorai (2015), “Inattention and Inertia in Household Finance: Evidence from the Danish Mortgage Market,” NBER Working Paper No. 21386.
- [2] Anscombe, F.J. and R. J. Aumann (1963), “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199-205.
- [3] De Oliveira, H., T. Denti, M. Mihm, and M.K. Ozbek (2016), “Rationally inattentive preferences,” *Theoretical Economics*, forthcoming.
- [4] Dillenberger, D., and P. Sadowski (2012), “Subjective Learning,” PIER Working Paper number 12-007.
- [5] Dillenberger, D., R.V. Krishna, P. Sadowski (2016), “Subjective Dynamic Information Constraints,” ERID working paper 214.
- [6] Dillenberger, D., J.S. Lleras, P. Sadowski, and N. Takeoka (2014), “A Theory of Subjective Learning,” *Journal of Economic Theory*, 153, 287–312.
- [7] Ghirardato, P. (2002), “Revisiting Savage in a conditional world,” *Economic Theory*, 20, 83–92.
- [8] Hodgkinson, G. (1997), “Cognitive inertia in a turbulent market: The case of U.K. residential estate agents,” *Journal of Management Studies*, 34: 921–945.
- [9] Kreps, D. (1979), “A Representation Theorem for “Preference for Flexibility,” *Econometrica*, 47, 565–577.

- [10] Lu, J. (2015), “Random Choice and Private Information,” mimeo.
- [11] Mackowiak, B. and M. Wiederholt (2009) “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99(3), 769-803.
- [12] Mackowiak, B. and M. Wiederholt (2015) “Business Cycle Dynamics under Rational Inattention,” *Review of Economic Studies*, 82(4), 1502–1532.
- [13] Samuelson, W., and R. Zeckhauser (1988), “Status Quo Bias in Decision Making,” *Journal of Risk and Uncertainty* 1, 7–59.
- [14] Savage, J. (1954), “The Foundations of Statistics.” New York: John Wiley & Sons.
- [15] Sims, C. A. (2003), “Implications of Rational Inattention”. *Journal of Monetary Economics*, 50 (3), 665–690.
- [16] Sims, C. A. (1998), “Stickiness”. Carnegie-Rochester Conference Series on Public Policy 49, 317–356.
- [17] Siniscalchi, M. (2011), “Dynamic Choice under Ambiguity,” *Theoretical Economics*, 6(3), 379–421.
- [18] Shapley, L. S. (1967), “On Balanced Sets and Cores.” *Naval Research Logistics Quarterly*, 14, 453–460.
- [19] Steiner, J., C. Stewart, and F. Matějka, (2015), “Rational Inattention Dynamics: Inertia and Delay in Decision-Making,” Tech. rep. CERGE-EI.
- [20] Tripsas, M., and G. Gavetti (2000), “Capabilities, Cognition, and Inertia: Evidence from Digital Imaging,” *Strategic Management Journal*, 21, 1147–1161.